RUGGEDCOM®
INDUSTRIAL STRENGTH NETWORKS™

Rugged Operating System

(ROS®) v3.7.0 User Guide

For use with :

RSG2200, M2200

July 16, 2009

Rugged Operating System: (ROS®) v3.7.0 User Guide

Copyright © 2009 RuggedCom Inc.

ALL RIGHTS RESERVED

Dissemination or reproduction of this document, or evaluation and communication of its contents, is not authorized except where expressly permitted. Violations are liable for damages. All rights reserved, particularly for the purposes of patent application or trademark registration.

This document contains proprietary information, which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of RuggedCom Inc.

Disclaimer Of Liability

We have checked the contents of this manual against the hardware and software described. However, deviations from the description cannot be completely ruled out.

RuggedCom shall not be liable for any errors or omissions contained herein or for consequential damages in connection with the furnishing, performance, or use of this material.

The information given in this document is reviewed regularly and any necessary corrections will be included in subsequent editions. We appreciate any suggested improvements. We reserve the right to make technical improvements without notice.

Registered Trademarks

RuggedServer™, RuggedWireless™, RuggedCom Discovery Protocol™ (RCDP™), RuggedExplorer™, Enhanced Rapid Spanning Tree Protocol™ (eRSTP™), are trademarks of RuggedCom Inc. Rugged Operating System® (ROS®) and RuggedSwitch® are registered trademarks of RuggedCom Inc. Other designations in this manual might be trademarks whose use by third parties for their own purposes would infringe the rights of the owner.

Warranty

Five (5) years from date of purchase, return to factory. For warranty details, visit www.ruggedcom.com or contact your customer service representative.

Contacting RuggedCom

Corporate Headquarters	US Headquarters	Europe Headquarters
RuggedCom Inc. 30 Whitmore Road Woodbridge, Ontario Canada, L4L 7Z4 Tel: (905) 856-5288 Fax: (905) 856-1995 Toll-free: 1 (888) 264-0006	RuggedCom 1930 Harrison St., Suite 209 Hollywood, Florida USA, 33020 Tel: (954) 922-7938x103 Fax: (954) 922-7984 Toll-free: 1 (888) 264-0006	RuggedCom Unit 41, Aztec Centre, Aztec West, Almondsbury, Bristol United Kingdom BS32 4TD Tel: +44 1454 203 404 Fax: +44 1454 203 403
	Email: RuggedSales@RuggedCom.com	

Technical Support
Toll Free (North America): 1 (866) 922-7975 International: +1 (905) 856-5288 Email: Support@RuggedCom.com

Web: www.RuggedCom.com

Table of Contents

Preface 12
Supported Platforms 12
Who Should Use This User Guide 12
How Chapters are organized 12
Document Conventions 12
Applicable Firmware Revision 13
Firmware/User Guide Version Numbering System 13
1. Administration 14
1.1. The ROS® User Interface 14
1.1.1. Using the RS232 Port to Access the User Interface 14
1.1.2. The Structure of the User Interface 14
1.1.3. Making Configuration Changes 15
1.1.4. Updates Occur In Real Time 16
1.1.5. Alarm Indications Are Provided 16
1.1.6. The CLI Shell 16
1.2. The ROS® Secure Shell Server 16
1.2.1. Using a Secure Shell to Access the User Interface 16
1.2.2. Using a Secure Shell to Transfer Files 16
1.3. The ROS® Web Server Interface 17
1.3.1. Using a Web Browser to Access the Web Interface 17
1.3.2. The Structure of the Web Interface 19
1.3.3. Making Configuration Changes 20
1.3.4. Updating Statistics Displays 21
1.4. Administration Menu 21
1.5. IP Interfaces 22
1.6. IP Gateways 25
1.7. IP Services 26
1.8. System Identification 27
1.9. Passwords 28
1.10. Time Synchronization 31
1.10.1. Configuring Time and Date 32
1.10.2. Configuring IRIG B 33
1.10.3. Time Source Selection 34
1.10.4. Time Sync Status 34
1.11. SNMP Management 35
1.11.1. SNMP Users 35
1.11.2. SNMP Security to Group Maps 37
1.11.3. SNMP Access 38
1.12. RADIUS 40
1.12.1. RADIUS overview 40
1.12.2. User Login Authentication and Authorization 40
1.12.3. 802.1X Authentication 41
1.12.4. Radius Server Configuration 42
1.13. TACACS+ 43
1.13.1. User Login Authentication and Authorization 43
1.13.2. TACACS+ Server Configuration 44

1.14. DHCP Relay Agent 45
1.15. Syslog 46
1.15.1. Configuring Local Syslog 47
1.15.2. Configuring Remote Syslog Client 47
1.15.3. Configuring the Remote Syslog Server 48
1.16. Troubleshooting 49
2. Ethernet Ports 50
2.1. Controller Protection Through Link-Fault-Indication (LFI) 50
2.2. Ethernet Ports Configuration and Status 51
2.2.1. Port Parameters 52
2.2.2. Port Rate Limiting 55
2.2.3. Port Mirroring 57
2.2.4. Cable Diagnostics 58
2.2.5. Link Detection Options 62
2.2.6. Port Status 63
2.2.7. Resetting Ports 64
2.3. Troubleshooting 64
3. Ethernet Statistics 66
3.1. Viewing Ethernet Statistics 66
3.2. Viewing Ethernet Port Statistics 68
3.3. Clearing Ethernet Port Statistics 72
3.4. Remote Monitoring (RMON) 72
3.4.1. RMON History Controls 73
3.4.2. RMON History Samples 74
3.4.3. RMON Alarms 78
3.5. RMON Events 82
3.6. RMON Event Log 84
4. Link Aggregation 86
4.1. Link Aggregation Operation 86
4.1.1. Link Aggregation Rules 86
4.1.2. Link Aggregation Limitations 88
4.2. Link Aggregation Configuration 89
4.2.1. Configuring Port Trunks 89
5. Spanning Tree 91
5.1. RSTP Operation 91
5.1.1. RSTP States and Roles 92
5.1.2. Edge Ports 94
5.1.3. Point-to-Point and Multipoint Links 94
5.1.4. Path and Port Costs 94
5.1.5. Bridge Diameter 95
5.2. MSTP Operation 95
5.2.1. MST Regions and Interoperability 96
5.2.2. MSTP Bridge and Port Roles 97
5.2.3. Benefits of MSTP 98
5.2.4. Implementing MSTP on a Bridged Network 99
5.3. RSTP Applications 100
5.3.1. RSTP in Structured Wiring Configurations 100
5.3.2. RSTP in Ring Backbone Configurations 101
5.3.3. RSTP Port Redundancy 103

5.4. Spanning Tree Configuration 103
5.4.1. Bridge RSTP Parameters 105
5.4.2. Port RSTP Parameters 107
5.4.3. MST Region Identifier 110
5.4.4. Bridge MSTI Parameters 111
5.4.5. Port MSTI Parameters 112
5.5. Spanning Tree Statistics 114
5.5.1. Bridge RSTP Statistics 114
5.5.2. Port RSTP Statistics 116
5.5.3. Bridge MSTI Statistics 118
5.5.4. Port MSTI Statistics 120
5.6. Troubleshooting 122
6. VLANs 125
6.1. VLAN Operation 125
6.1.1. VLANs and Tags 125
6.1.2. Tagged vs. Untagged Frames 125
6.1.3. Native VLAN 125
6.1.4. Management VLAN 126
6.1.5. Edge and Trunk Port Types 126
6.1.6. VLAN Ingress and Egress Rules 127
6.1.7. Forbidden Ports List 127
6.1.8. VLAN-aware and VLAN-unaware Operation Modes 127
6.1.9. GVRP (GARP VLAN Registration Protocol) 128
6.1.10. PVLAN Edge 129
6.1.11. QinQ 130
6.2. VLAN Applications 131
6.2.1. Traffic Domain Isolation 131
6.2.2. Administrative Convenience 132
6.2.3. Reduced Hardware 132
6.3. VLAN Configuration 132
6.3.1. Global VLAN Parameters 133
6.3.2. Static VLANs 134
6.3.3. Port VLAN Parameters 135
6.3.4. VLAN Summary 137
6.4. Troubleshooting 138
7. Port Security 140
7.1. Port Security Operation 140
7.1.1. Static MAC address-based authorization 140
7.1.2. IEEE 802.1X Authentication 140
7.2. Port Security Configuration 142
7.2.1. Port Security Parameters 142
7.2.2. 802.1X Parameters 144
8. Classes of Service 146
8.1. CoS Operation 146
8.1.1. Inspection Phase 146
8.1.2. Forwarding Phase 147
8.2. CoS Configuration 147
8.2.1. Global CoS Parameters 148
8.2.2. Port CoS Parameters 149

8.2.3. Priority to CoS Mapping 150
8.2.4. DSCP to CoS Mapping 152
9. Multicast Filtering 154
9.1. IGMP 154
9.1.1. Router and Host IGMP Operation 154
9.1.2. Switch IGMP Operation 155
9.1.3. Combined Router and Switch IGMP Operation 157
9.2. GMRP (GARP Multicast Registration Protocol) 158
9.2.1. GMRP Example 159
9.3. Multicast Filtering Configuration and Status 161
9.3.1. Configuring IGMP Parameters 162
9.3.2. Configuring Static Multicast Groups 164
9.3.3. Viewing IP Multicast Groups 165
9.3.4. Configuring GMRP 166
9.4. Troubleshooting 167
10. MAC Address Tables 170
10.1. Viewing MAC Addresses 171
10.2. Configuring MAC Address Learning Options 172
10.3. Configuring Static MAC Address Table 172
10.4. Purging MAC Address Table 174
11. Network Discovery 175
11.1. LLDP Operation 175
11.2. RCDP Operation 176
11.3. Network Discovery Menu 176
11.3.1. LLDP Menu 177
11.3.2. RCDP Configuration 184
12. Diagnostics 185
12.1. Using the Alarm System 185
12.1.1. Active Alarms 185
12.1.2. Passive Alarms 186
12.1.3. Alarms and the Critical Failure Relay 186
12.1.4. Configuring Alarms 186
12.1.5. Viewing and Clearing Alarms 189
12.2. Viewing CPU Diagnostics 190
12.3. Viewing and Clearing the System Log 191
12.4. Viewing Product Information 192
12.5. Loading Factory Default Configuration 193
12.6. Resetting the Device 193
13. Using the CLI Shell 195
13.1. Summary Of CLI Commands available in ROS® 195
13.2. Obtaining Help For A Command 196
13.3. Viewing Files 196
13.3.1. Listing Files 196
13.3.2. Viewing and Clearing Log Files 197
13.4. Pinging a Remote Device 197
13.5. Tracing Events 198
13.5.1. Enabling Trace 199
13.5.2. Starting Trace 199
13.6. Viewing DHCP Learned Information 200

13.7. Executing Commands Remotely Through RSH 200
13.8. Resetting the Device 201
14. Firmware Upgrade and Configuration Management 202
14.1. Upgrading Firmware 202
14.1.1. Upgrading Firmware Using XModem 203
14.1.2. Upgrading Firmware Using the ROS TFTP Server 204
14.1.3. Upgrading Firmware Using the ROS® TFTP Client 204
14.1.4. Upgrading Firmware Using SFTP 205
14.2. Updating Configuration 205
14.3. Backing Up ROS System Files 206
14.3.1. Backing Up Files Using SFTP 207
14.4. Using SQL Commands 207
14.4.1. Getting Started 207
14.4.2. Finding the Correct Table 208
14.4.3. Retrieving Information 208
14.4.4. Changing Values in a Table 209
14.4.5. Setting Default Values in a Table 209
14.4.6. Using RSH and SQL 209
A. SNMP MIB Support 211
A.1. Standard MIBs 211
A.2. RuggedCom proprietary MIBs 212
B. SNMP Trap Summary 213
C. List of Objects Eligible for RMON Alarms 214
D. ModBus Management Support and Memory Map 219
D.1. Modbus Memory Map 220
D.1.1. Text 223
D.1.2. Cmd 223
D.1.3. Uint16 224
D.1.4. Uint32 224
D.1.5. PortCmd 224
D.1.6. Alarm 225
D.1.7. PSStatusCmd 225
E. Command Line Listing 227
Index 230

List of Figures

1.1. Main Menu With Screen Elements Identified 15
1.2. Log in to The Device with a Web Browser 18
1.3. Log In To The Device With A Web Browser (Custom Login Banner; banner.txt Empty) 19
1.4. Main Menu via Web Server Interface 20
1.5. Web Page Header Showing Alarms Link 20
1.6. Parameters Form Example 21
1.7. Administration Menu 22
1.8. IP Interfaces Table 23
1.9. IP Interfaces Form 24
1.10. IP Gateways Form 25
1.11. IP Services Form 26
1.12. System Identification Form 28
1.13. Passwords Form 29
1.14. Time Synchronization Menu 31
1.15. Time and Date Form 32
1.16. Serial Port Menu 33
1.17. Time Source Selection Menu 34
1.18. Time Sync Status Menu 34
1.19. SNMP User Table 35
1.20. SNMP User Form 36
1.21. SNMP Security to Group Maps Table 37
1.22. SNMP Security to Group Maps Form 37
1.23. SNMP Access Table 38
1.24. SNMP Access Form 39
1.25. RADIUS Server Summary 42
1.26. RADIUS Server Form 42
1.27. TACACS+ Server Summary 44
1.28. TACACS+ Server Form 44
1.29. DHCP Relay Agent Form 46
1.30. Local Syslog Form 47
1.31. Remote Syslog Client Form 47
1.32. Remote Syslog Server Table 48
1.33. Remote Syslog Server Form 48
1.34. Using A Router As A Gateway 49
2.1. Controller Protection Through LFI 50
2.2. Ethernet Ports Menu 52
2.3. Port Parameters Table 52
2.4. Port Parameters Form 53
2.5. Port Rate Limiting Table 55
2.6. Port Rate Limiting Form 56
2.7. Port Mirroring Form 57
2.8. Cable Diagnostics Table 59
2.9. Cable Diagnostics Parameters Form 59
2.10. Link Detection Form 62
2.11. Port Status Table 63

3.1. Ethernet Port Statistics Menu 66
3.2. Ethernet Statistics Table 67
3.3. Ethernet Port Statistics Table 68
3.4. Ethernet Port Statistics Form 69
3.5. Clear Ethernet Port Statistics Form 72
3.6. RMON History Controls Table 73
3.7. RMON History Controls Form 73
3.8. RMON History Samples Table 75
3.9. RMON History Samples Form 76
3.10. The Alarm Process 79
3.11. RMON Alarms Table 79
3.12. RMON Alarms Form 80
3.13. RMON Events Table 82
3.14. RMON Events Form 83
3.15. RMON Event Log Table 84
3.16. RMON Event Log Form 84
4.1. Link Aggregation Examples 86
4.2. Displaying Port Trunk Secondary Ports in Layer 2 Feature Configuration 87
4.3. Link Aggregation Menu 89
4.4. Port Trunk Table 89
4.5. Port Trunk Form 90
5.1. Bridge and Port States 92
5.2. Bridge and Port Roles 93
5.3. Example of a Structured Wiring Configuration 100
5.4. Example of a Ring Backbone Configuration 102
5.5. Port Redundancy 103
5.6. Spanning Tree Menu 104
5.7. Bridge RSTP Parameters Form 105
5.8. Port RSTP Parameter Table 107
5.9. Port RSTP Parameter Form 108
5.10. MST Region Identifier Form 110
5.11. Bridge MSTI Parameters 111
5.12. Port MSTI Parameter Table 112
5.13. Port MSTI Parameter Form 112
5.14. Bridge RSTP Statistics Form 114
5.15. Port RSTP Statistics Table 116
5.16. Port RSTP Statistics Form 116
5.17. Bridge MSTI Statistics Form 118
5.18. Port MSTI Statistics Table 120
5.19. Port MSTI Statistics Form 120
6.1. Using GVRP 129
6.2. Using QinQ Example 130
6.3. Multiple Overlapping VLANs 131
6.4. Inter-VLAN Communications 132
6.5. Virtual LANs Menu 133
6.6. Global VLAN Parameters Form 133
6.7. Static VLANs Table 134
6.8. Static VLANs Form 134
6.9. Port VLAN Parameters Table 135

6.10. Port VLAN Parameters Form 136
6.11. VLAN Summary Table 138
7.1. 802.1X General Topology 141
7.2. 802.1X Packet Exchange 141
7.3. Port Security Menu 142
7.4. Port Security Parameters Table 142
7.5. Port Security Parameters Form 143
7.6. 802.1X Parameters Table 144
7.7. 802.1X Parameters Form 144
8.1. Determining The CoS Of A Received Frame 147
8.2. Classes Of Service Menu 148
8.3. Global CoS Parameters Form 148
8.4. Port CoS Parameter Form 149
8.5. Port CoS Parameter Form 149
8.6. Priority to CoS Mapping Table 150
8.7. Priority to CoS Mapping Form 151
8.8. TOS DSCP to CoS Mapping Table 152
8.9. TOS DSCP to CoS Mapping Form 152
9.1. IGMP Operation Example 1 155
9.2. IGMP Operation Example 2 157
9.3. Example using GMRP 160
9.4. Multicast Filtering Menu 162
9.5. IGMP Parameters Form 162
9.6. Static Multicast Groups Table 164
9.7. Static Multicast Group Form 164
9.8. IP Multicast Groups Table 165
9.9. GMRP Port Summary 166
9.10. Port GMRP Parameters 166
9.11. Multicast Group Summary 167
10.1. MAC Address Tables Menu 170
10.2. Address Table 171
10.3. MAC Address Learning Options Form 172
10.4. Static MAC Address Table 173
10.5. Static MAC Address Form 173
11.1. Network Discovery Main Menu 177
11.2. Network Discovery Menu 178
11.3. Global LLDP Parameters Form 179
11.4. Port LLDP Parameters Table 180
11.5. Port LLDP Parameters Form 180
11.6. LLDP Global Remote Statistics Form 181
11.7. LLDP Neighbor Information Table 182
11.8. LLDP Statistics Table 183
11.9. LLDP Neighbor Information Table 184
12.1. Diagnostics Menu 185
12.2. Alarm Configuration Table 187
12.3. Alarm Configuration Form 188
12.4. Alarm Table 189
12.5. CPU Diagnostics Form 190
12.6. Viewing the System Log 191

12.7. Product Information Form 192
12.8. Load Factory Defaults Dialog 193
12.9. Reset Device Dialog 194
13.1. Displaying The List Of Available Commands 196
13.2. Displaying Help For A Command 196
13.3. Displaying The Directory Of A ROS Device 197
13.4. Displaying Trace Settings 198
13.5. Enabling Trace 199
13.6. Starting Trace 200

Preface

This manual contains instructions, examples, guidelines, and general theory on how to use the Rugged Operating System (ROS®) management software.

Supported Platforms

ROS has been designed to work on many RuggedCom product hardware platforms. This ensures consistency of the user experience when migrating from one product model to another. In fact, a single 'binary' image supports all RuggedCom ROS-based products that includes:

- RuggedSwitch® i800, i801, i802, and i803
- RuggedSwitch® RS8000 and RS1600
- RuggedSwitch® RS900/RS930 with both 'L' (EoVDSL) and 'W' (WLAN) port variants
- RuggedSwitch® RS900G/RS940G with Gigabit
- RuggedSwitch® RS969/M969 waterproof with Gigabit
- RuggedSwitch® RSG2100/M2100 and RSG2200/M2200 modular switches with Gigabit Ethernet
- RuggedSwitch® RSG2300 modular switch with Gigabit Ethernet
- RuggedServer™ RS416, RS910 and RS920 modular servers
- RuggedServer™ RS400
- RuggedServer™ RMC30

Each product model has a subset of the entire ROS feature set. This manual is intended for use with the RSG2200, M2200 product family and has been streamlined to only describe the relevant features.

Who Should Use This User Guide

This guide is to be used by network technical support personnel who are familiar with the operation of networks. Others who might find the book useful are network and system planners, system programmers and line technicians.

How Chapters are organized

The index of this guide has been prepared with:

- Entries to each of the "Features" sections of the manual
- Entries to each of the "Troubleshooting" sections of the manual (located at the end of each chapter)
- Entries to each of the Menus, organized by name

Document Conventions

This publication uses the following conventions:

Note

Means reader take note. Notes contain helpful suggestions or references to materials not contained in this guide.

It is recommended that you use this guide along with the following applicable documents:

- RSG2200, M2200 Family Installation Guide
- RuggedCom Fiber Guide
- RuggedCom Wireless Guide
- White paper: Rapid Spanning Tree in Industrial Networks

Applicable Firmware Revision

This guide is applicable to ROS software revisions v3.7.x.

Firmware/User Guide Version Numbering System

ROS has a three-digit version numbering system of the form X.Y.Z where each digit is a number starting from zero. The 'X.Y' digits represent the functional version of ROS whereas the 'Z' digit represents firmware patches. The 'X' digit is incremented for major functional updates of the product. The 'Y' digit is incremented for minor functional updates of the product. The 'Z' digit is incremented for bug fixes, cosmetic enhancements and other minor issues.

User guides follow the same format. In general, a user guide will have the same 'X.Y' digits as the firmware to which it corresponds.

It is RuggedCom's policy to provide Web access to only the latest 'patch' release for a version of firmware. If you decide that an upgrade is merited, then getting all the fixes only makes sense. It is for this reason that release notes are created detailing all patches for a given functional version.

1. Administration

The Administration menu covers the configuration of administrative parameters of both device and network (local services availability, security methods employed, system identification and functionality related to the IP network):

- IP Address, Subnet Mask and Gateway Address (static or dynamically obtainable)
- Management VLAN
- Management Connection Inactivity Timeout
- TFTP Server Permissions
- System Identification
- Passwords
- Time and Date
- SNTP to keep the time and date synchronized
- SNMP Management
- Radius Server
- DHCP Relay Agent
- Remote Syslog

1.1. The ROS® User Interface

1.1.1. Using the RS232 Port to Access the User Interface

Attach a terminal (or PC running terminal emulation software) to the RS232 port. The terminal should be configured for 8 bits, no parity operation at 57.6 Kbps. Hardware and software flow control must be disabled. Select a terminal type of VT100.

Once the terminal is connected, pressing any key on the keyboard will prompt for the user name and password to be entered.

The switch is shipped with a default administrator user name - “admin” - and password - “admin”. Once successfully logged in, the user will be presented with the main menu.

1.1.2. The Structure of the User Interface

The user interface is organized as a series of menus with an escape to a command line interface (CLI) shell. Each menu screen presents the switch name (as provided by the System Identification parameter), Menu Title, Access Level, Alarms indicator, Sub-Menus and Command Bar.

Sub-menus are entered by selecting the desired menu with the arrow keys and pressing the enter key. Pressing the escape key returns you to the parent menu.

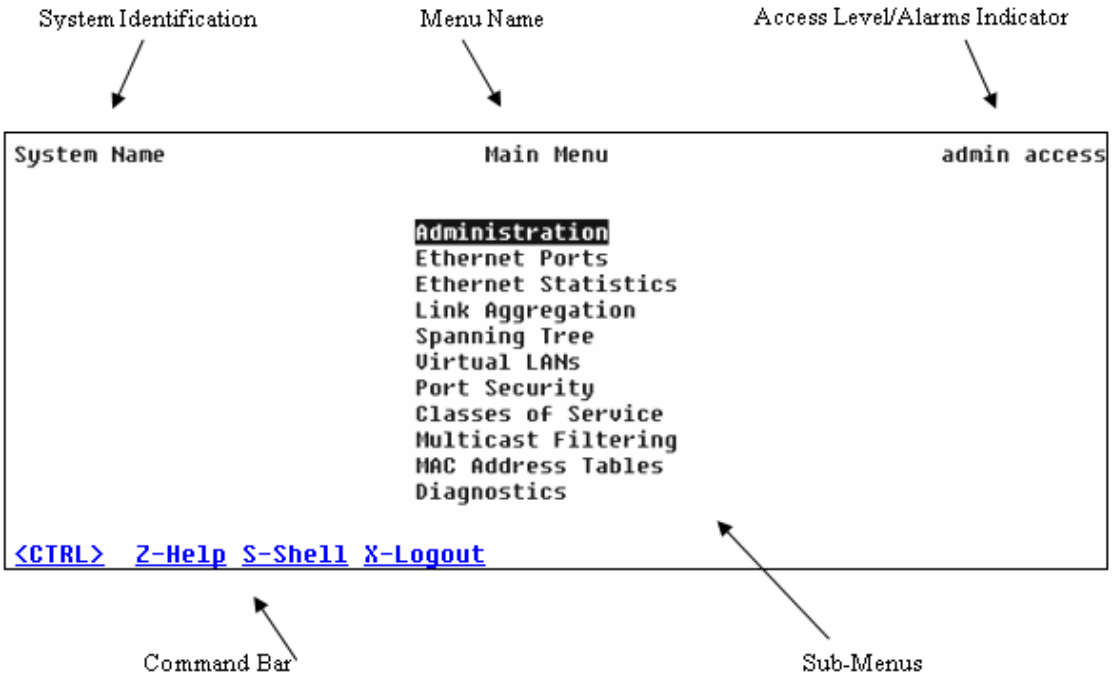

Figure 1.1. Main Menu With Screen Elements Identified

The command bar offers a list of commands that apply to the currently displayed menu. These commands include:

- <Ctrl-Z> to display help on the current command or data item
- <Ctrl-S> to switch to the CLI shell
- <Ctrl-Up/Down> to jump to next/previous page of a status display

The main menu also provides a <Ctrl-X> command, which will terminate the session. This type of menu is accessible via serial console, telnet session and SSH session.

1.1.3. Making Configuration Changes

When changing a data item, the user selects the data item by the cursor keys and then pressing the enter key. The cursor will change position to allow editing of the data item.

Typing a new value after pressing enter always erases the old parameter value. The left and right cursor keys can be used to position the edit point without erasing the old parameter value. The up and down cursor keys can be used to cycle through the next higher and lower values for the parameter.

After the parameter has been edited, press enter again to change other parameters. When all desired parameters have been modified, press <Ctrl-A> to apply changes. The switch will automatically prompt you to save changes when you leave a menu in which changes have been made.

Some menus will require you to press <Ctrl-I> to insert a new record of information and <Ctrl-L> to delete a record.

1.1.4. Updates Occur In Real Time

All configuration and display menus present the current values, automatically updating if changed from other user interface sessions or SNMP. All statistics menus will display changes to statistics as they occur.

1.1.5. Alarm Indications Are Provided

Alarms are events for which the user is notified through the Diagnostics sub-menu. All configuration and display menus present an indication of the number of alarms (in the upper right hand corner of the screen) as they occur, automatically updating as alarms are posted and cleared.

1.1.6. The CLI Shell

The user interface provides a Command Line Interface shell for operations that are more easily performed at the command line. You may switch back and forth from the menu system and shell by pressing <Ctrl-S>. For more information on the capabilities of the shell please refer to Chapter 13, *Using the CLI Shell*.

1.2. The ROS® Secure Shell Server

1.2.1. Using a Secure Shell to Access the User Interface

SSH (Secure Shell) is a network protocol which provides a replacement for insecure remote login and command execution facilities, such as telnet and remote shell. SSH encrypts traffic in both directions, preventing traffic sniffing and password theft.

SSH protocol version 2 is implemented in ROS. The authentication method is "keyboard-interactive" password authentication. A user logged in via SSH has the same privileges as one logged in via the console port.

1.2.2. Using a Secure Shell to Transfer Files

ROS implements an SFTP server via SSH to transfer files securely. The file system visible on the RuggedSwitch® has a single directory. The files in it are created at startup time and can be neither deleted nor renamed. Existing files can be downloaded from the switch. For example, firmware images may be downloaded for backup and log files may be downloaded for analysis. Some files may be overwritten by uploading a file of the same name to the switch, as would be done in order to upgrade the firmware.

The implemented commands are:

dir/ls
: list directory contents

get
: download a file from the switch

put
: upload a file to the switch

The files that may be overwritten via SFTP upload are:

main.bin
main ROS firmware image

boot.bin
RuggedSwitch bootloader image

config.csv
ROS configuration file

fpga.xsvf
FPGA configuration file

1.3. The ROS® Web Server Interface

1.3.1. Using a Web Browser to Access the Web Interface

A web browser uses a secure communications method called SSL (Secure Socket Layer) to encrypt traffic exchanged with its clients. The web server guarantees that communications with the client are kept private. If the client requests access via an insecure HTTP port, it will be rerouted to the secure port. Access to the web server via SSL will be granted to a client that provides a valid user name / password pair.

Note

It can happen that upon connecting to the ROS web server, a web browser may report that it cannot verify the authenticity of the server's certificate against any of its known certificate authorities. This is expected, and it is safe to instruct the browser to accept the certificate. Once the browser accepts the certificate, all communications with the web server will be secure.

Start a web browser session and open a connection to the switch by entering a URL that specifies its host name or IP address (e.g. http://179.1.0.45). Once in contact with the switch, start the login process by clicking on the “Login” link. The resulting page should be similar to that presented below:

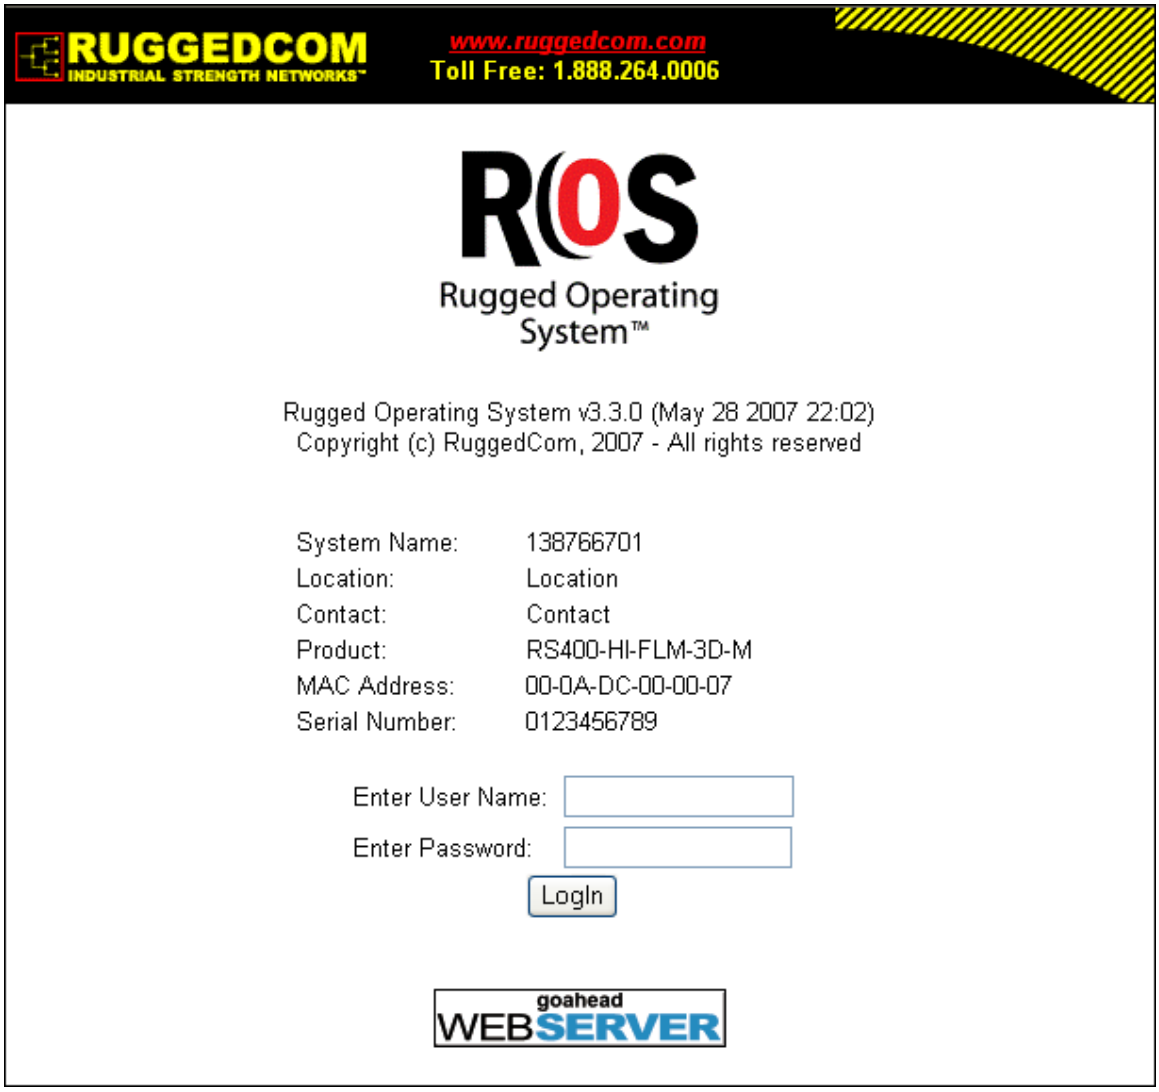

Figure 1.2. Log in to The Device with a Web Browser

Enter the “admin” user name and the appropriate password for the admin user, and then click on the “LogIn” button. The switch is shipped with a default administrator password of “admin”. Once successfully logged in, the user will be presented with the main menu.

If the user wants to hide the device information from the login screen, the ‘Login Banner’ option in the System Identification menu must be set to ‘Custom’, and the desired data to be displayed should be uploaded to the device as the file: "banner.txt".

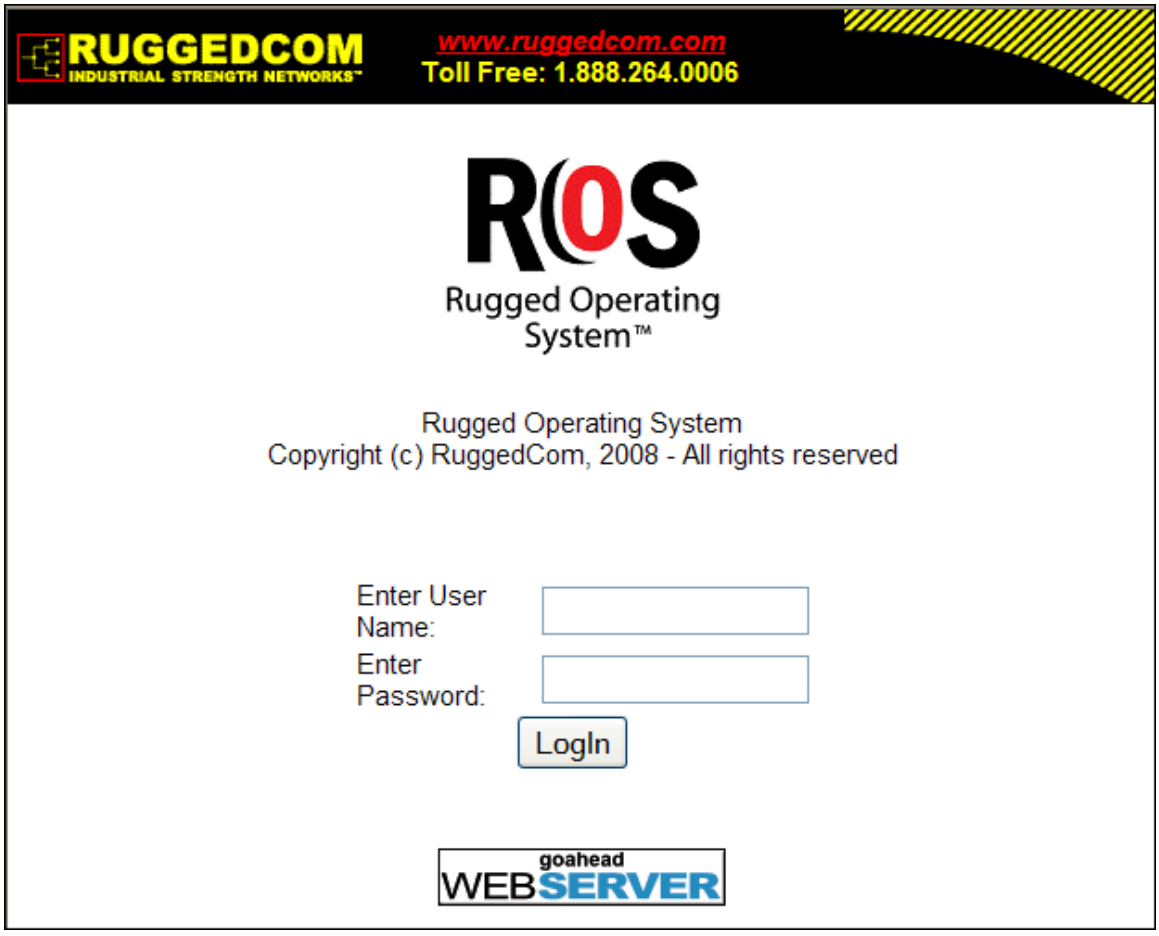

Figure 1.3. Log In To The Device With A Web Browser (Custom Login Banner; banner.txt Empty)

1.3.2. The Structure of the Web Interface

The user interface is organized as a series of linked web pages. The main menu provides the links at the top level of the menu hierarchy and allows them to be expanded to display lower-level links for each configuration subsystem.

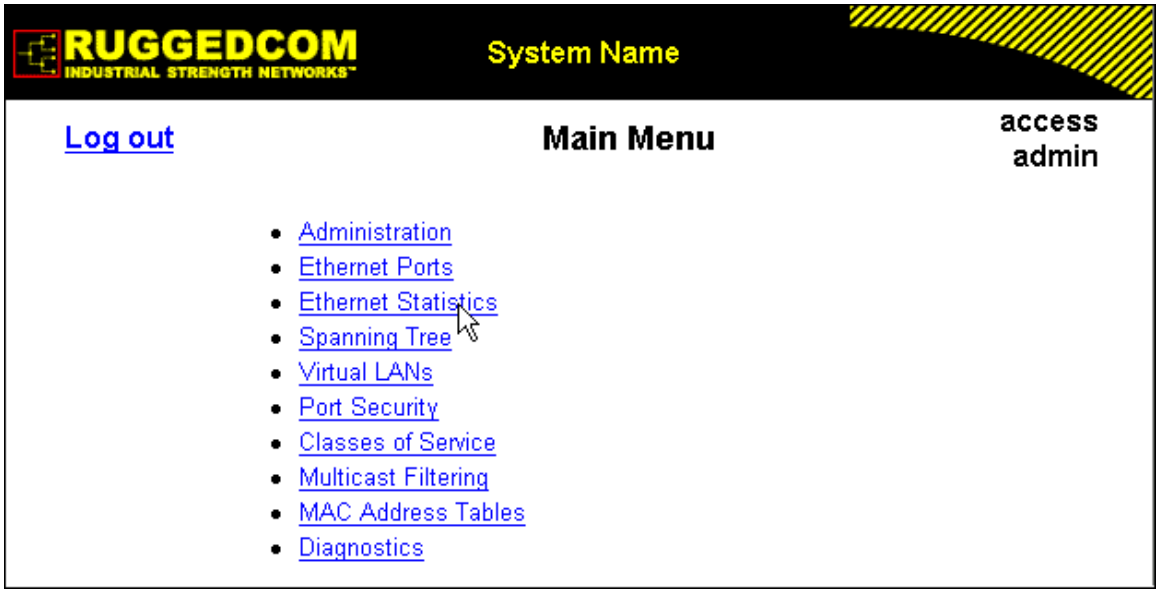

Figure 1.4. Main Menu via Web Server Interface

Every web page in the menu system has a common header section which contains:

- The System Name, as configured in the System Identification menu, is displayed in the top banner, in between elements of the RuggedCom logo.
- A "Log out" link at left and immediately below the banner, terminates the current web session.
- A "Back" link at left and below "Log out" links back to the previously viewed page.
- The menu title, in the center of the page and below the banner, is a link to a context-sensitive help page.
- The access level, e.g. "access admin", is displayed by default at the right of the page and below the banner. If, however, any alarms are pending, the text will be replaced with a link which displays the number of pending alarms. Following this link displays a table of pending alarms.

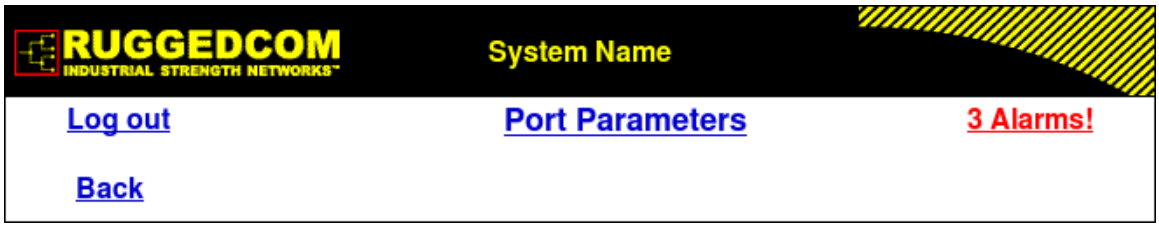

Figure 1.5. Web Page Header Showing Alarms Link

1.3.3. Making Configuration Changes

When changing a data item, the user selects the data item by selecting the field to edit with the mouse, entering a new value and clicking on the apply field. More than one parameter may be modified at a time.

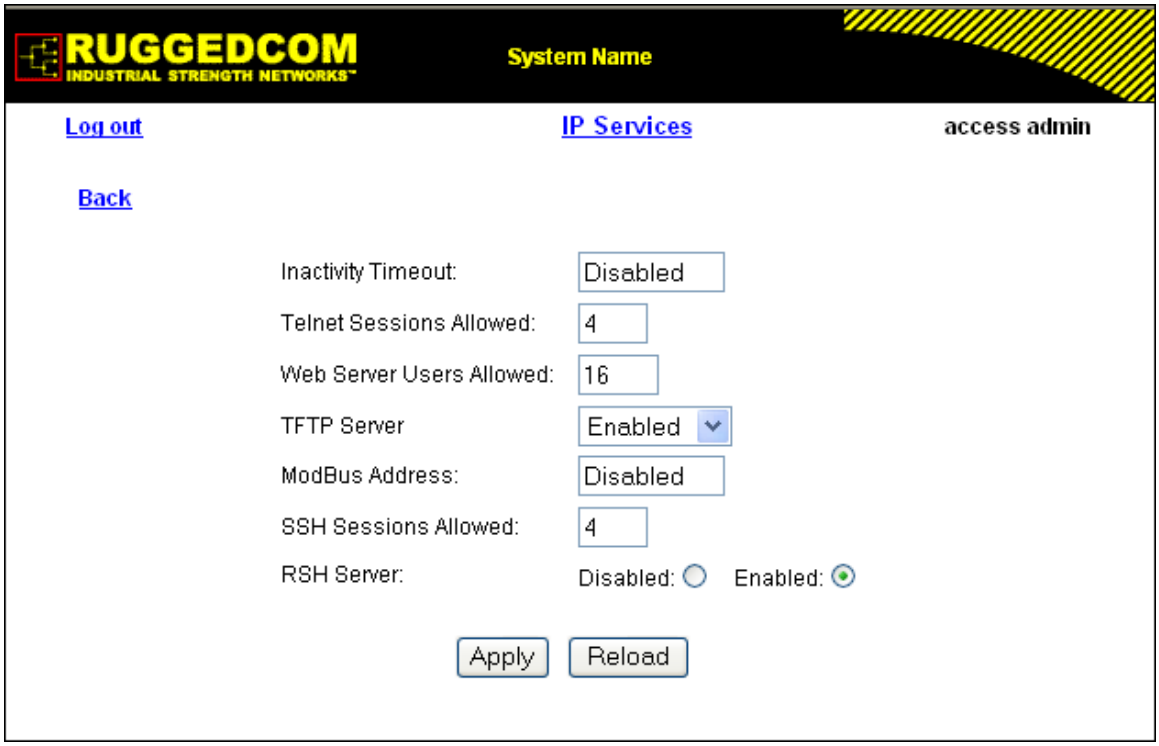

Figure 1.6. Parameters Form Example

Some menus will require you to create or delete new records of information.

1.3.4. Updating Statistics Displays

You may click the refresh button to update statistics displays.

1.4. Administration Menu

The Administration menu provides ability to configure network and switch administration parameters.

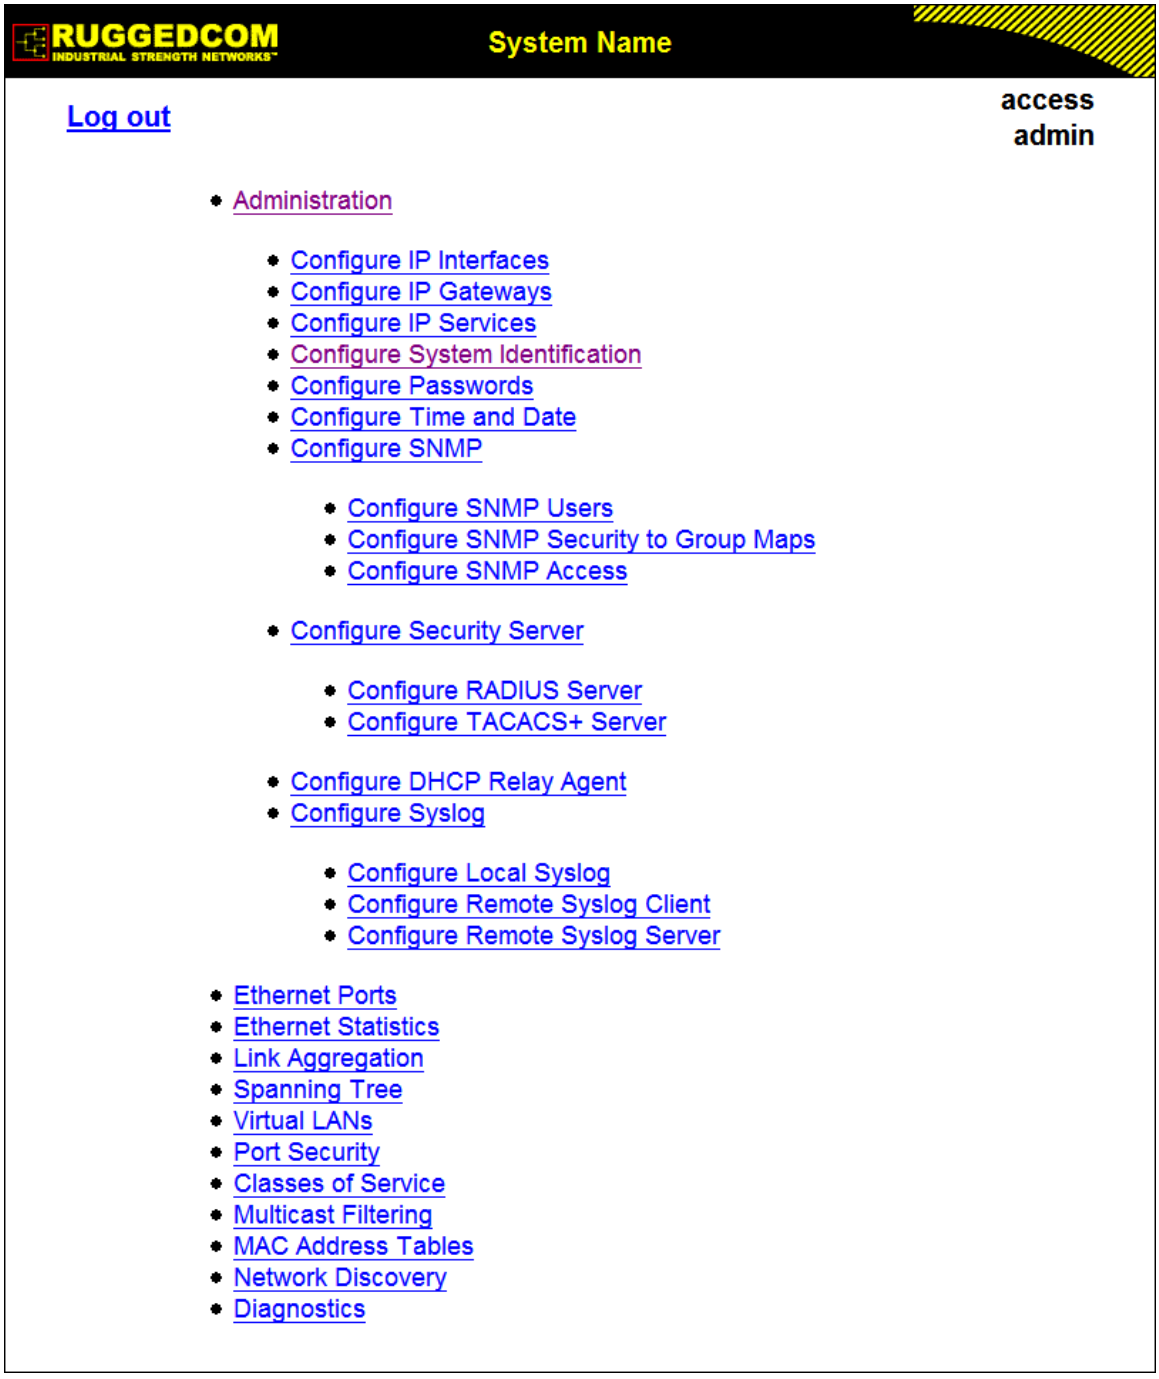

Figure 1.7. Administration Menu

1.5. IP Interfaces

These parameters provide the ability to configure IP connection parameters such as address, network, and mask.

The user can configure an IP interface for each subnet (VLAN). One of the interfaces is configured to be the management interface. IP services: TFTP server, SNMP server, Telnet server, SSH server, RSH server, Web server, authentication using a RADIUS server, DHCP client, BOOTP

client or DHCP relay agent. A DHCP relay agent will be available only via the management interface. Different IP interfaces must not overlap, i.e. the subnet mask must be unique.

15 IP interfaces can be configured in the device. In VLAN unaware mode, and in devices that do not act as switches (as RMC30), only one IP interface can be configured.

On non-management interfaces, only static IP addresses can be assigned.

On the management interface, the user can choose from the following IP Address types: Static, DHCP, BOOTP and Dynamic. Static IP Address type refers to the manual assignment of an IP address while DHCP, BOOTP and Dynamic IP Address types refer to the automatic assignment of an IP address.

DHCP is widely used in LAN environments to dynamically assign IP addresses from a centralized server, which reduces the overhead of administrating IP addresses.

BOOTP is a subset of the DHCP protocol. ROS supports the transfer of a BOOTFILE via BOOTP. The BOOTFILE represents any valid ROS file such as config.csv. The name of BOOTFILE on the BOOTP server must match the corresponding ROS file.

The Dynamic IP Address type refers to a combination of the BOOTP and DHCP protocols. Starting with BOOTP, the system will try BOOTP and DHCP in a round-robin fashion until it receives a response from the corresponding server.

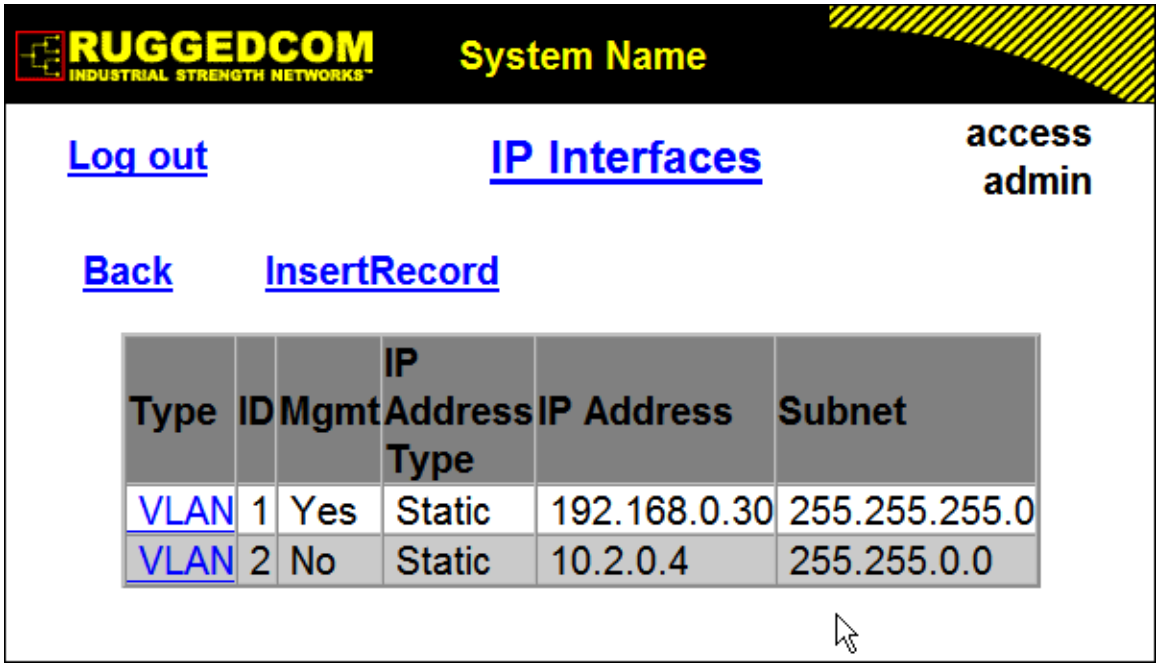

Figure 1.8. IP Interfaces Table

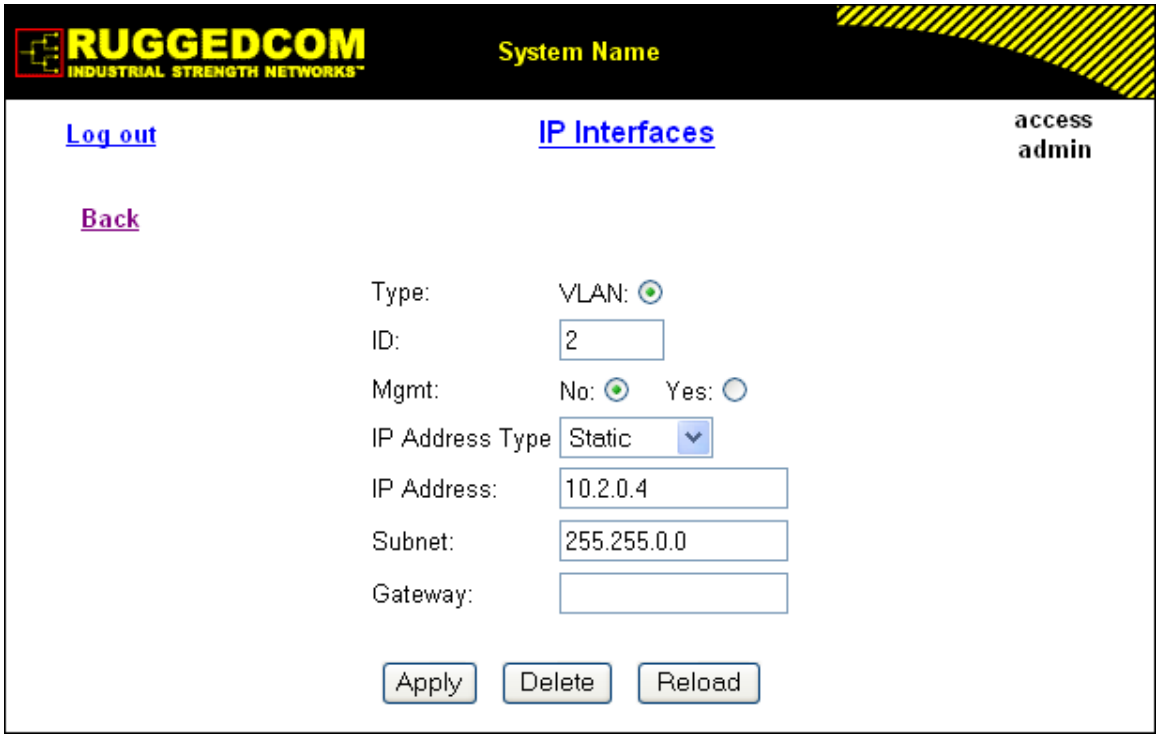

Figure 1.9. IP Interfaces Form

Note

The IP address and mask configured for the management VLAN are not changed when resetting all configuration parameters to defaults and will be assigned a default VLAN ID of 1. Changes to the IP address take effect immediately. All IP connections in place at the time of an IP address change will be lost.

Type

```
Synopsis: { VLAN }
Default: VLAN
```

Specifies the type of the interface for which this IP interface is created.

ID

```
Synopsis: 1 to 4094
Default: 1
```

Specifies the ID of the interface for which this IP interface is created. If the interface type is VLAN, this represents the VLAN ID.

Mgmt

```
Synopsis: { No, Yes }
Default: No
```

Specifies whether the IP interface is the device management interface.

IP Address Type

```
Synopsis: { Static, Dynamic, DHCP, BOOTP }
Default: Static
```

Specifies whether the IP address is static or is dynamically assigned via DHCP or BOOTP. The Dynamic option automatically switches between BOOTP and DHCP until it receives a response from the relevant server. The Static option must be used for non-management interfaces.

IP Address

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default: 192.168.0.1
```

Specifies the IP address of this device. An IP address is a 32-bit number that is notated by using four numbers from 0 through 255, separated by periods. Only a unicast IP address is allowed, which ranges from 1.0.0.0 to 233.255.255.255.

Subnet

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default: 255.255.255.0
```

Specifies the IP subnet mask of this device. An IP subnet mask is a 32-bit number that is notated by using four numbers from 0 through 255, separated by periods. Typically, subnet mask numbers use either 0 or 255 as values (e.g. 255.255.255.0) but other numbers can appear.

1.6. IP Gateways

These parameters provide the ability to configure gateways. A maximum of 10 gateways can be configured. When both the Destination and Subnet fields are both 0.0.0.0 (displayed as blank space), the gateway is a default gateway.

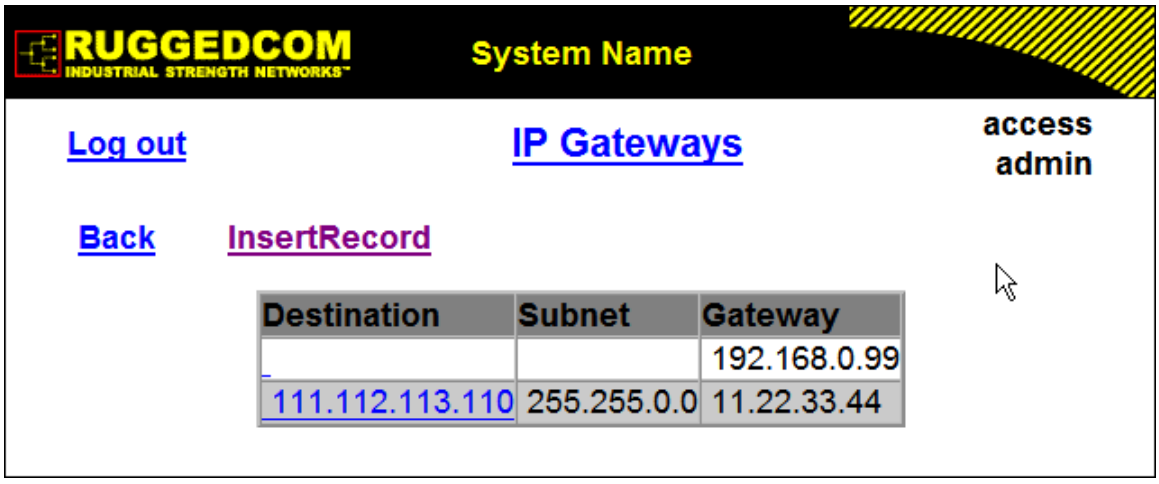

Figure 1.10. IP Gateways Form

Destination

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default: 0.0.0.0
```

Specifies the IP address of the destination device. An IP address is a 32-bit number that is notated by using four numbers from 0 through 255, separated by periods.

Subnet

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
```

```
Default: 0.0.0.0
```

Specifies the IP subnet mask of the destination. An IP subnet mask is a 32-bit number that is notated by using four numbers from 0 through 255, separated by periods. Typically, subnet mask numbers use either 0 or 255 as values (e.g. 255.255.255.0) but other numbers can appear.

Gateway

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default: 0.0.0.0
```

Specifies the gateway IP address. The gateway address must be on the same IP subnet as this device.

Note

The default gateway configuration will not be changed when resetting all configuration parameters to defaults.

1.7. IP Services

These parameters provide the ability to configure properties for IP services provided by the device.

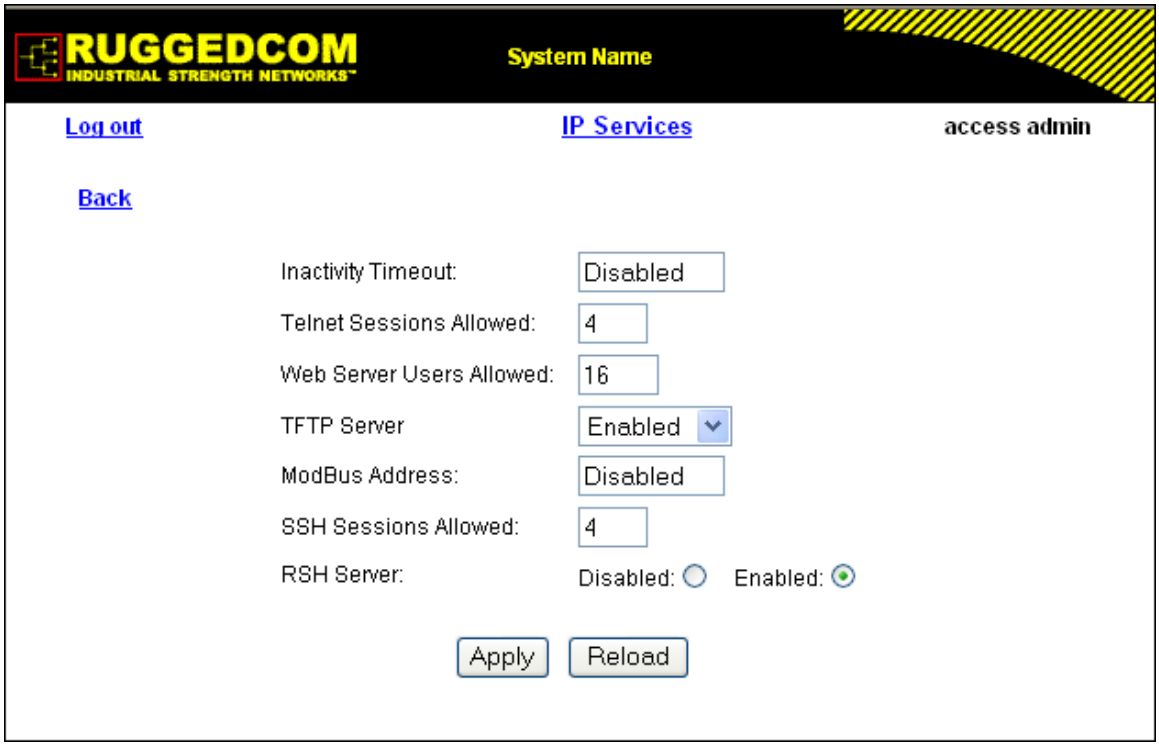

Figure 1.11. IP Services Form

Inactivity Timeout

```
Synopsis: 1 to 60 or { Disabled }
Default: 5 min
```

Specifies when the console will timeout and display the login screen if there is no user activity. A value of zero disables timeouts for console and Telnet users. For Web Server users maximum timeout value is limited to 30 minutes.

Telnet Sessions Allowed

```
Synopsis: 0 to 4
Default: 4
```

Limits the number of Telnet sessions. A value of zero prevents any Telnet access.

Web Server Users Allowed

```
Synopsis: 1 to 16
Default: 16
```

Limits the number of simultaneous web server users.

TFTP Server

```
Synopsis: { Disabled, Get Only, Enabled }
Default: Get Only
```

As TFTP is a very insecure protocol, this parameter allows the user to limit or disable TFTP Server access.

DISABLED - disables read and write access to TFTP Server

GET ONLY - only allows reading of files via TFTP Server

ENABLED - allows reading and writing of files via TFTP Server

ModBus Address

```
Synopsis: 1 to 254 or { Disabled }
Default: Disabled
```

Determines the Modbus address to be used for Management through Modbus.

SSH Sessions Allowed

```
Synopsis: 1 to 4
Default: 4
```

Limits the number of SSH sessions.

RSH Server

```
Synopsis: { Disabled, Enabled }
Default: Enabled
```

Disables/enables Remote Shell access.

1.8. System Identification

The system identification is displayed in the sign-on screen and in the upper left hand corner of all ROS screens.

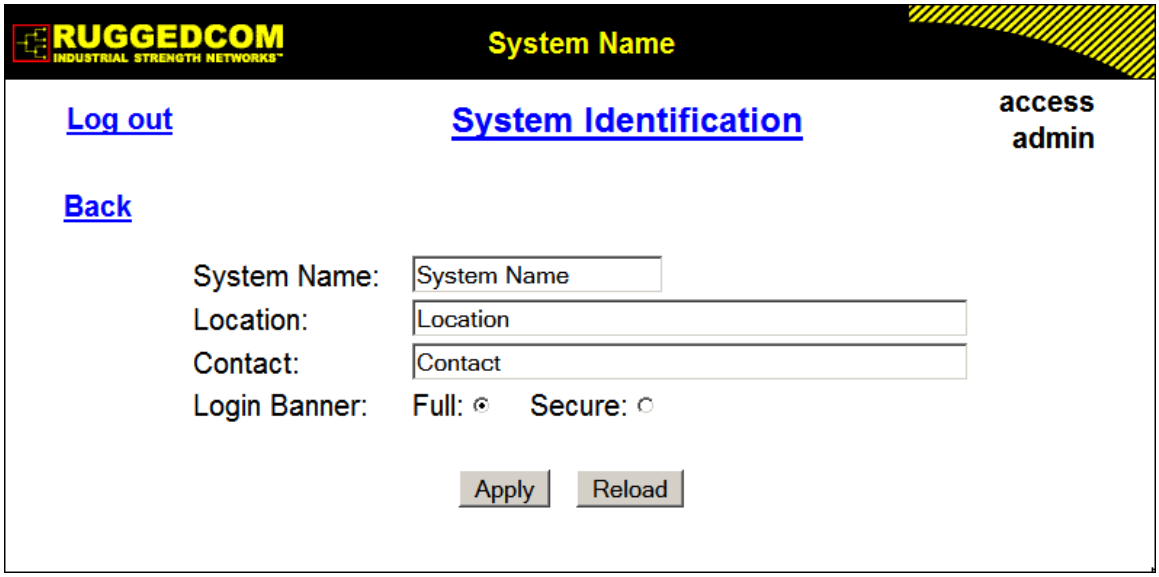

Figure 1.12. System Identification Form

System Name

```
Synopsis: Any 19 characters
Default: System Name
```

The system name is displayed in all ROS menu screens. This can make it easier to identify the switches within your network, provided that all switches are given a unique name.

Location

```
Synopsis: Any 49 characters
Default: Location
```

The location can be used to indicate the physical location of the switch. It is displayed in the login screen as another means to ensure you are dealing with the desired switch.

Contact

```
Synopsis: Any 49 characters
Default: Contact
```

The contact can be used to help identify the person responsible for managing the switch. You can enter name, phone number, email, etc. It is displayed in the login screen so that this person may be contacted, should help be required.

Login Banner

```
Synopsis: { Standard, Custom }
Default: Standard
```

Provides the ability to customize the banner displayed on the login screen. Either the standard RuggedCom ROS banner may be displayed or, if "Custom" is selected, the contents of a file named "banner.txt", uploaded to the device, will be used as a login banner.

1.9. Passwords

These parameters provide the ability to configure parameters for authorized and authenticated access to the device's services (HMI via Serial Console, Telnet, SSH, RSH, Web Server). Access

to the switch can be authorized and authenticated via RADIUS or TACACS+ servers, or using locally configured passwords that are configured per user name and access level.

Note that access via the Serial Console is always authorized first using local settings. If a local match is not found, RADIUS/TACACS+ will be used if enabled. For all other services, if RADIUS/TACACS+ is enabled for authentication and authorization, the local setting will be used only if configured.

To access the unit, the user name and password must be provided.

Three user names and passwords can be configured. They correspond to three access levels, which provide or restrict access to change settings and execute various commands within the device.

- **guest** users can view most settings, but may not change settings or run commands.
- **operator** cannot change settings, but can reset alarms, clear statistics and logs.
- **admin** user can change all the settings and run commands.

Figure 1.13. Passwords Form

Auth Type

```
Synopsis: { Local, RADIUS, TACACS+, RADIUSorLocal, TACACS+orLocal }
Default: Local
```

Password authentication can be performed using locally configured values, a remote RADIUS server, or a remote TACACS+ server. Setting this value to one of the combinations that includes RADIUS or TACACS+ requires that the Security Server Table be configured.

- Local - authentication from the local Password Table

- RADIUS - authentication using a RADIUS server
- TACACS+ - authentication using a TACACS+ server
- RADIUSOrLocal - authentication using RADIUS. If the server cannot be reached, authenticate from the local Password Table.
- TACACS+OrLocal - authentication using TACACS+. If the server cannot be reached, authenticate from the local Password Table

Guest Username

```
Synopsis: 15 character ASCII string
Default: guest
```

Related password is in the Guest Password field; view only, cannot change settings or run any commands.

Guest Password

```
Synopsis: 15 character ASCII string
Default: guest
```

Related user name is in the Guest Username field; view only, cannot change settings or run any commands.

Confirm Guest Password

```
Synopsis: 15 character ASCII string
Default: None
```

Confirm the input of the above Guest Password.

Operator Username

```
Synopsis: 15 character ASCII string
Default: operator
```

Related password is in the Oper Password field; cannot change settings; can reset alarms, statistics, logs, etc.

Operator Password

```
Synopsis: 15 character ASCII string
Default: operator
```

Related user name is in the Oper Username field; cannot change settings; can reset alarms, statistics, logs, etc.

Confirm Operator Password

```
Synopsis: 15 character ASCII string
Default: None
```

Confirm the input of the above Operator Password.

Admin Username

```
Synopsis: 15 character ASCII string
Default: admin
```

Related password is in the Admin Password field; full read/write access to all settings and commands.

Admin Password

```
Synopsis: 15 character ASCII string
Default: admin
```

Related user name is in the Admin Username field; full read/write access to all settings and commands.

Confirm Admin Password

```
Synopsis: 15 character ASCII string
Default: None
```

Confirm the input of the above Admin Password.

1.10. Time Synchronization

ROS supports:

- Local hardware time keeping and time zone management
- IRIG B input and output
- Generic PPS (Pulse Per Second) output
- SNTP time synchronization

In addition to the local clock and existing SNTP client/server, IRIG B may be configured to serve as the unit's master clock source.

The **Configure Time Sync** option within the ROS Administration menu fully configures time keeping functions on a ROS-based device:

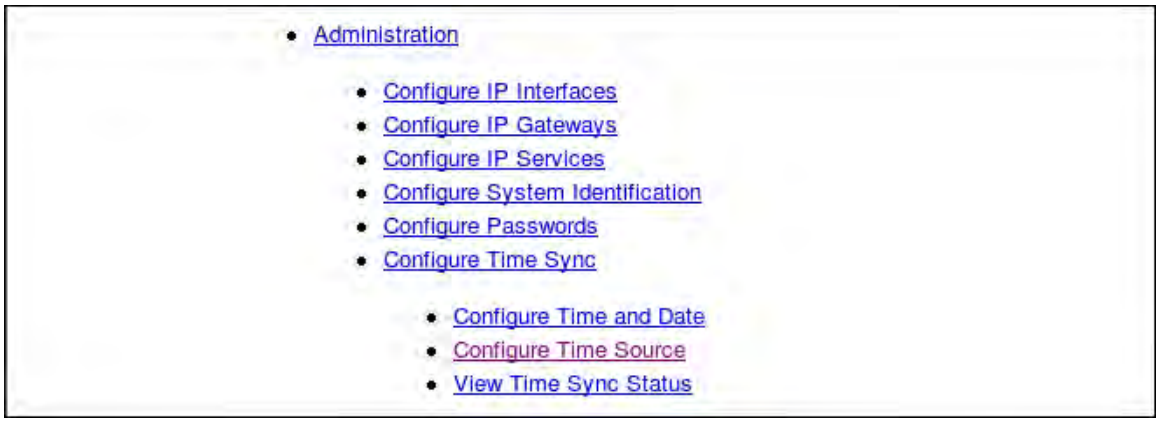

Figure 1.14. Time Synchronization Menu

IRIG B Fundamentals

The Inter-Range Instrumentation Group (IRIG) IRIG B standard details the format of a signal encoding which contains the current day, hour, minute and second in UTC format, broadcast at the start of each second. The IRIG B daughter card, available in the RS416v2, complies with IRIG Standard 200-04, and is capable of generating formats IRIGB002 and IRIGB003 (PWM).

Generic PPS (Pulse Per Second) Interface

The serial IRIG B daughter card also be configured to provide a generic PPS signal on the serial connectors on the RS416v2. A pulse with a duration of 1 millisecond is output every second at the beginning of the second.

1.10.1. Configuring Time and Date

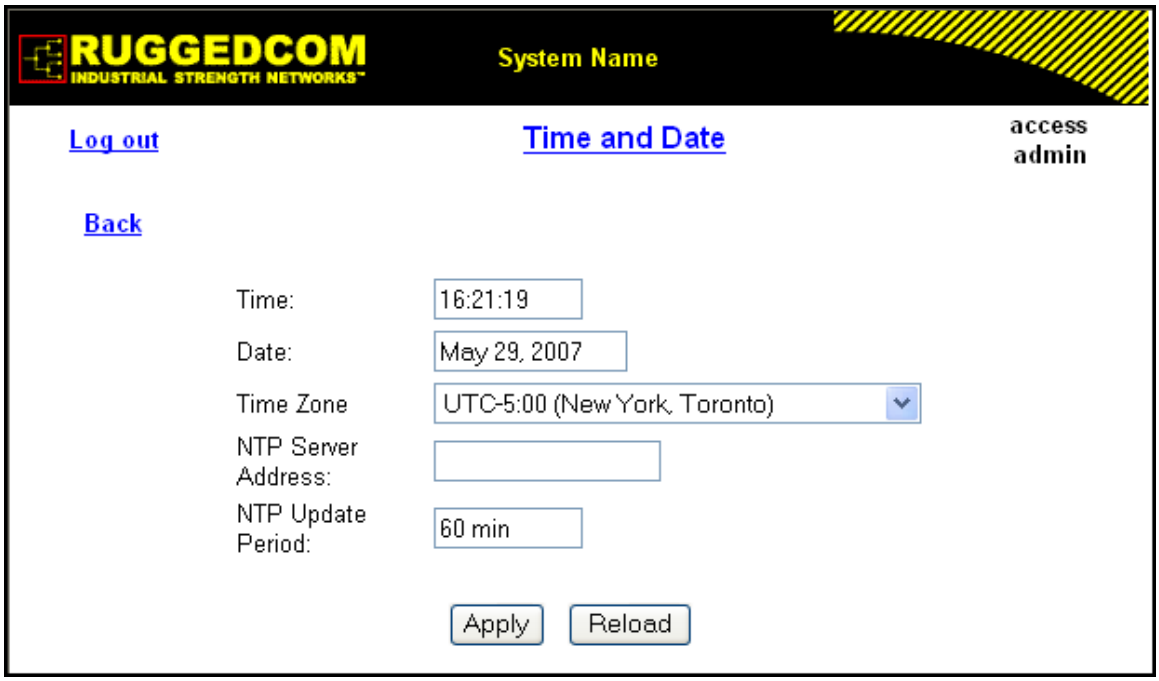

Figure 1.15. Time and Date Form

Device time, date and time zone can be set via this form. The device can also be configured to periodically contact an (S)NTP server to correct any drift in the onboard clock.

Each ROS-based device unit can act as a unicast SNTP server and/or SNTP client. The SNTP server will respond to the unicast SNTP requests received from the units where its address is configured as NTP Server Address. The server itself can be synchronized with higher-level NTP server.

Note

*If it is desired that NTP provide the time reference for the unit, the **Time Source** parameter must be set to "NTP Server" in the Time Source Selection menu.*

Time

```
Synopsis: HH:MM:SS
```

This parameter enables both the viewing and setting of the local time.

Date

```
Synopsis: MMM DD, YYYY
```

This parameter enables both the viewing and setting of the local date.

Time Zone

```
Synopsis: {
UTC-12:00 (Eniwetok, Kwajalein), UTC-11:00 (Midway Island, Samoa),
UTC-10:00 (Hawaii), UTC-9:00 (Alaska), UTC-8:00 (Los Angelos,
Vancouver),
UTC-7:00 (Calgary, Denver), UTC-6:00 (Chicago, Mexico City),
```

```
UTC-5:00 (New York, Toronto), UTC-4:00 (Caracas, Santiago),
UTC-3:30 (Newfoundland), UTC-3:00 (Brasilia, Buenos Aires),
UTC-2:00 (Mid Atlantic), UTC-1:00 (Azores),
UTC-0:00 (Lisbon, London), UTC+1:00 (Berlin, Paris, Rome),
UTC+2:00 (Athens, Cairo, Helsinki), UTC+3:00 (Baghdad, Moscow),
UTC+3:30 (Teheran), UTC+4:00 (Abu Dhabi, Kazan, Muscat),
UTC+4:30 (Kabul), UTC+5:00 (Islamabad, Karachi),
UTC+5:30 (Calcutta, New Delhi), UTC+5:45 (Kathmandu),
UTC+6:00 (Almaty, Dhaka), UTC+6:30 (Rangoon),
UTC+7:00 (Bangkok, Hanoi), UTC+8:00 (Beijing, Hong Kong)
UTC+9:00 (Seoul, Tokyo), UTC+9:30 (Adelaide, Darwin),
UTC+10:00 (Melbourne, Sydney), UTC+11:00 (Magadan, New Caledonia),
UTC+12:00 (Auckland, Fiji) }
Default: UTC-0:00 (Lisbon, London)
```

This setting enables the conversion of UTC (Universal Coordinated Time) to local time.

NTP Server Address

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default:
```

This parameter specifies the IP address of an (S)NTP server ((Simple) Network Time Protocol); programming an address of '0.0.0.0' disables SNTP requests. This device is an SNTP client which may connect to only one server. If a server address is programmed then a manual setting of the time will be overwritten at the next update period.

NTP Update Period

```
Synopsis: 1 to 1440
Default: 60 min
```

This setting determines how frequently the (S)NTP server is polled for a time update. If the server cannot be reached, three attempts are made at one-minute intervals and then an alarm is generated, at which point the programmed rate is resumed.

1.10.2. Configuring IRIG B

IRIG B functionality is configured via the serial port setup menu, since it is implemented on the same connectors.

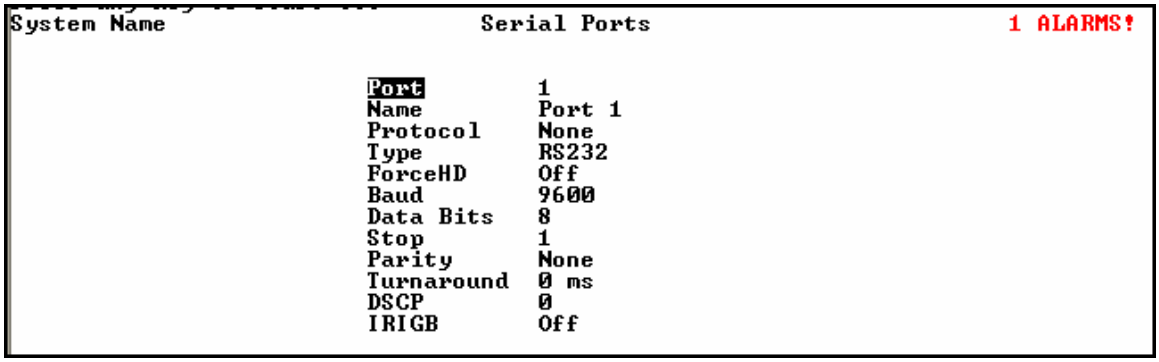

Figure 1.16. Serial Port Menu

A new field, **IRIGB** has been added to the **Serial Ports** configuration menu, which sets the operational mode of IRIGB port:

IRIGB

Synopsis: { PWM, PPS, Off }
Default: Off

Selects the output mode of the IRIG B pin on the corresponding serial port.

PWM (Pulse Width Modulation) mode complies with IRIG Standard 200-04 generating formats IRIGB002 and IRIGB003. In PPS mode, a pulse with a duration of 1 millisecond is output every second at the beginning of the second.

IRIG B as Reference Clock

IRIG B can be configured to act as a reference clock for the PTP and NTP system. The following are required to activate this functionality:

- IRIG-B daughter card is installed (presently only RS416 supports this card)
- Select IRIGB as a time source in the “Time Source” UI.

1.10.3. Time Source Selection

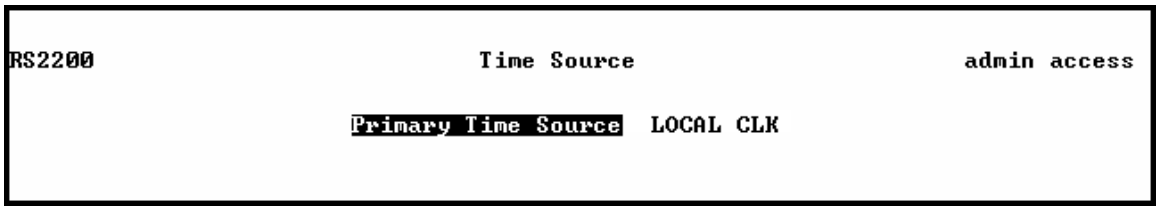

Figure 1.17. Time Source Selection Menu

Primary Time Source

Synopsis: { LOCAL CLK, IRIGB, NTP Server }
Default: NTP Server

Selects the time source that will drive the local clock and provide the reference for time protocols for which the device can act as a server. Note that changing the time source may produce a step change in the time seen via any of the clock outputs.

1.10.4. Time Sync Status

This menu provides summary information on the status of the time synchronization subsystem.

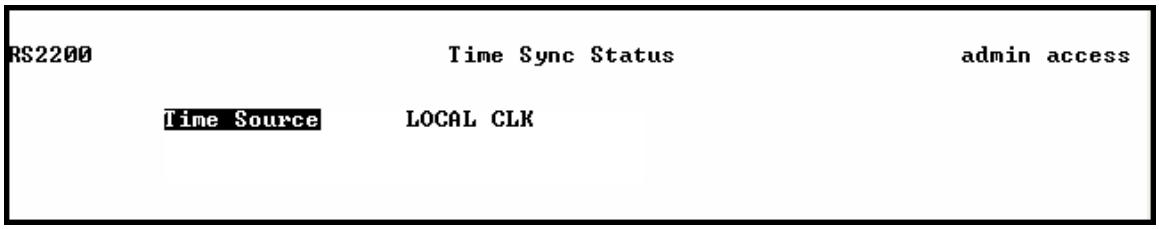

Figure 1.18. Time Sync Status Menu

Time Source

Displays the currently selected time source. Possible values are: IRIGB, NTP Server, and LOCAL CLK (local clock).

IRIGB Status

Displays the status of the IRIG B clock source. Note that this field is not displayed at all if IRIG B input hardware is not installed in the device.

1.11. SNMP Management

ROS supports Simple Network Management Protocol Version 3 (SNMPv3). This protocol provides secure access to devices by a combination of authentication and encrypting packets over the network. The security features provided are:

- message integrity - ensuring that a packet has not been tampered with in-transit.
- authentication – determining the message is from a valid source.
- encryption – scrambling the contents of a packet to prevent it from being seen by an unauthorized source.

SNMPv3 provides security models and security levels. A security model is an authentication strategy that is set up for a user and the group in which the user resides. A security level is a permitted level of security within a security model. A combination of a security model and security level will determine which security mechanism is employed when handling an SNMP packet.

Note the following about SNMPv3 protocol:

- each user belongs to a group.
- a group defines the access policy for a set of users.
- an access policy defines what SNMP objects can be accessed for: reading, writing and creating notifications.
- a group determines the list of notifications its users can receive.
- a group also defines the security model and security level for its users.

1.11.1. SNMP Users

These parameters provide the ability to configure users for the local SNMPv3 engine. Note that, if the security level employed is SNMPv1 or SNMPv2, User Name represents a community name for authentication or sending traps. Up to 32 entries can be configured.

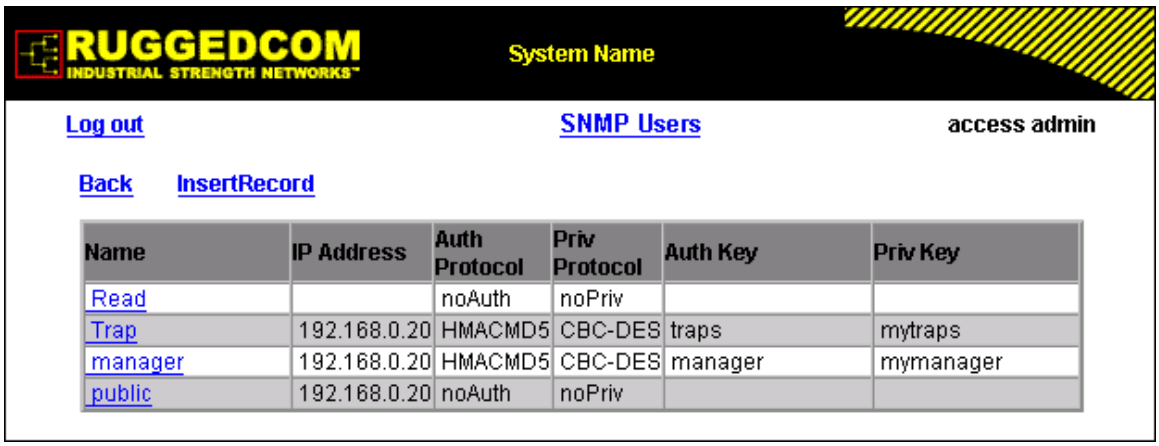

Figure 1.19. SNMP User Table

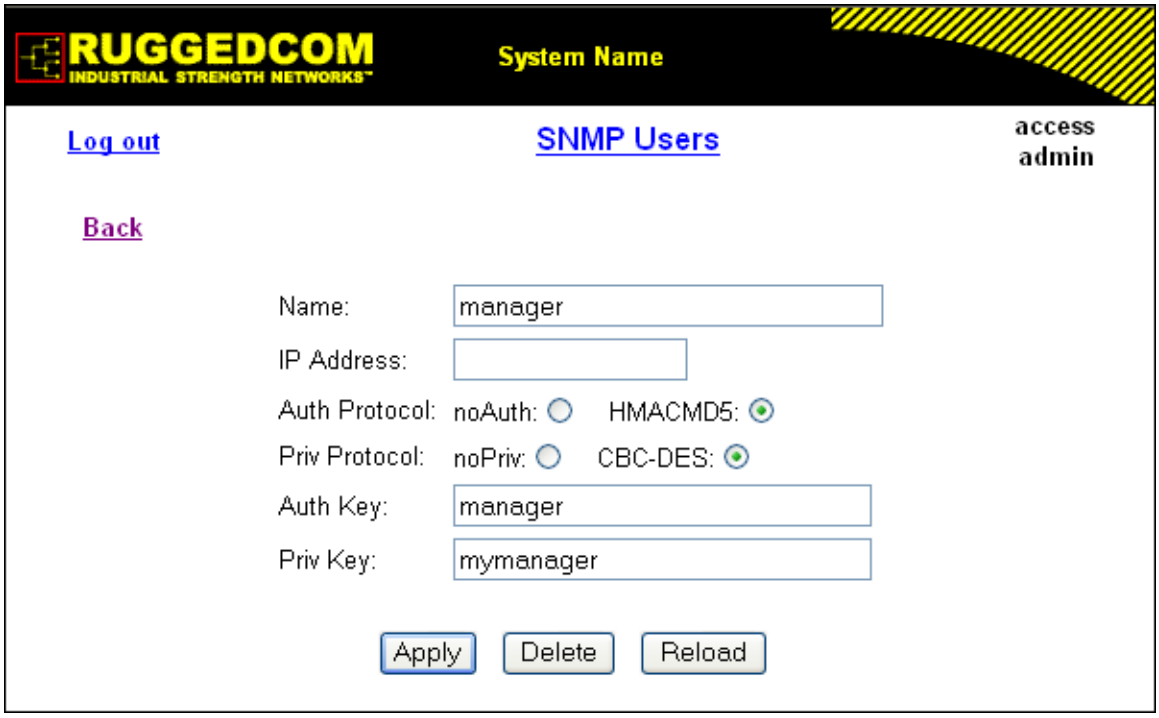

Figure 1.20. SNMP User Form

Name

```
Synopsis: Any 32 characters
Default: initial
```

The name of the user. This is the User-based Security Model dependent security ID.

IP Address

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default:
```

The IP address of the user's SNMP management station if it is configured to receive traps and notifications.

Auth Protocol

```
Synopsis: { noAuth, HMACMD5 }
Default: noAuth
```

An indication of whether messages sent on behalf of this user to/from the SNMP engine can be authenticated, and if so, the type of authentication protocol which is used.

Priv Protocol

```
Synopsis: { noPriv, CBC-DES }
Default: noPriv
```

An indication of whether the messages sent on behalf of this user to/from the SNMP engine can be protected from disclosure, and if so, the type of privacy protocol which is used.

Auth Key

```
Synopsis: 31 character ASCII string
Default: None
```

The secret authentication key (password) that must be shared with an SNMP client.

Confirm Auth Key

```
Synopsis: 31 character ASCII string
Default: None
```

Confirm input of the Auth Key, above.

Priv Key

```
Synopsis: 31 character ASCII string
Default: None
```

The secret encryption key (password) that must be shared with an SNMP client.

Confirm Priv Key

```
Synopsis: 31 character ASCII string
Default: None
```

Confirm input of the Priv Key, above.

1.11.2. SNMP Security to Group Maps

Entries in this table map configuration of security model and security name (user) into a group name, which is used to define an access control policy. Up to 32 entries can be configured.

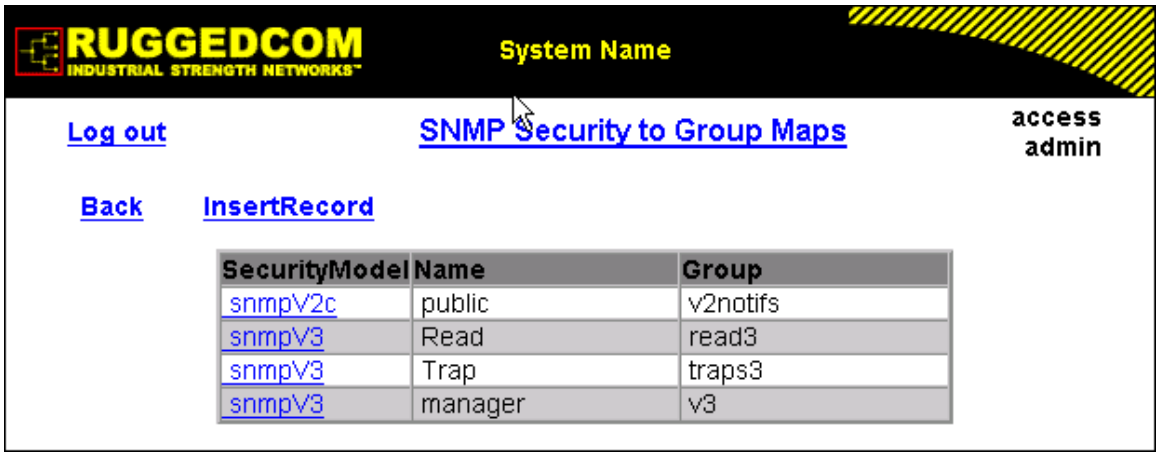

Figure 1.21. SNMP Security to Group Maps Table

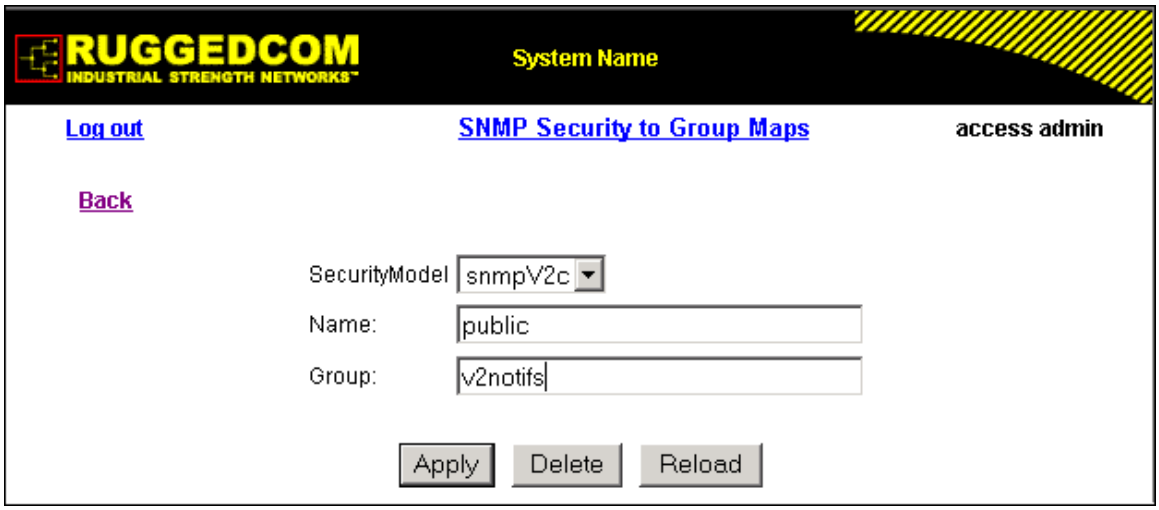

Figure 1.22. SNMP Security to Group Maps Form

SecurityModel

```
Synopsis: { snmpV1, snmpV2c, snmpV3 }
Default: snmpV3
```

The Security Model that provides the name referenced in this table.

Name

```
Synopsis: Any 32 characters
Default:
```

The user name which is mapped by this entry to the specified group name.

Group

```
Synopsis: Any 32 characters
Default:
```

The group name to which the security model and name belong. This name is used as an index to the SNMPv3 VACM Access Table.

1.11.3. SNMP Access

These parameters provide the ability to configure access rights for groups.To determine whether access is allowed, one entry from this table needs to be selected and the proper view name from that entry must be used for access control checking. View names are predefined:

- noView - access is not allowed
- V1Mib - SNMPv3 MIBs excluded
- allOfMibs - all supported MIBs are included.

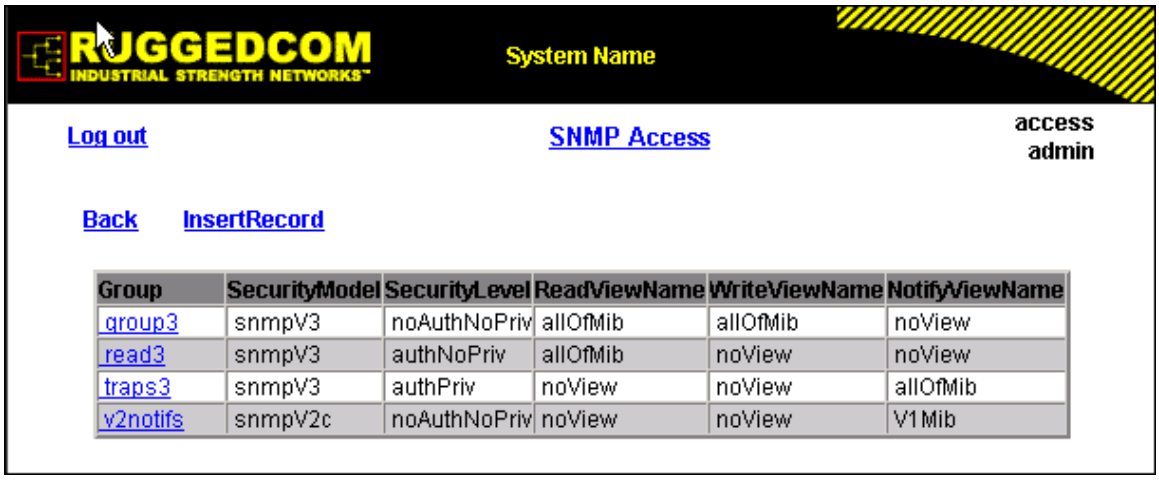

Figure 1.23. SNMP Access Table

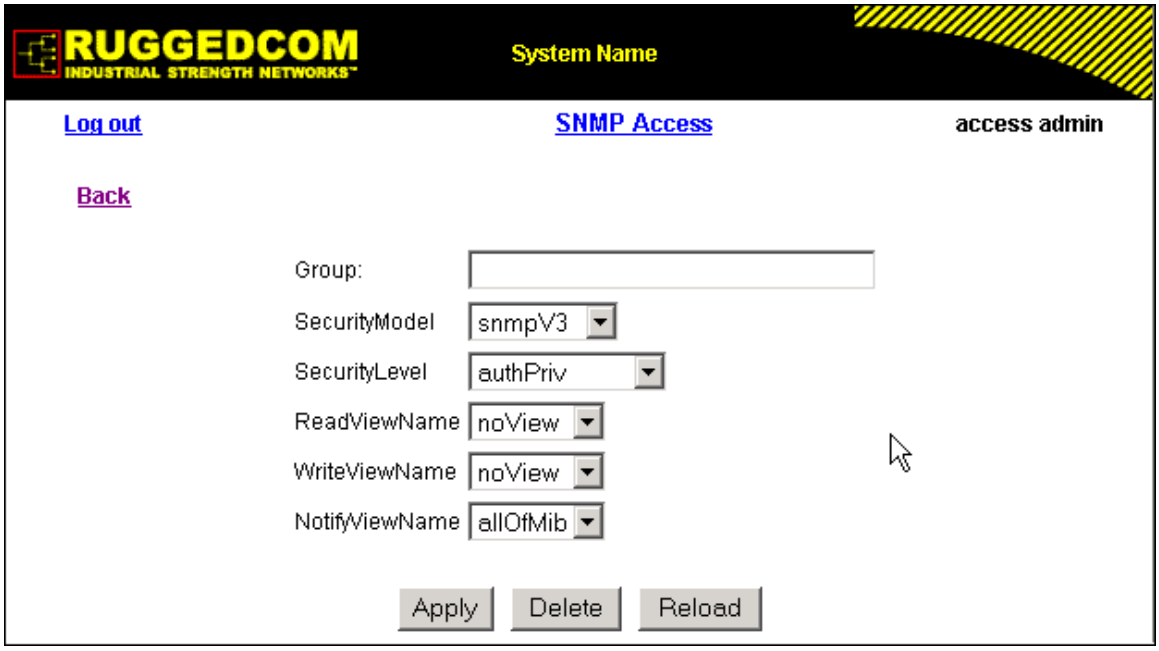

Figure 1.24. SNMP Access Form

Group

```
Synopsis: Any 32 characters
Default:
```

The group name to which the security model and name belong. This name is used as an index to the SNMPv3 VACM Access Table.

SecurityModel

```
Synopsis: { snmpV1, snmpV2c, snmpV3 }
Default: snmpV3
```

In order to gain the access rights allowed by this entry, the configured security model must be in use.

SecurityLevel

```
Synopsis: { noAuthNoPriv, authNoPriv, authPriv }
Default: noAuthNoPriv
```

The minimum level of security required in order to gain the access rights allowed by this entry. A security level of noAuthNoPriv is less than authNoPriv, which is less than authPriv.

ReadViewName

```
Synopsis: { noView, V1Mib, allOfMib }
Default: noView
```

This parameter identifies the MIB tree(s) to which this entry authorizes read access. If the value is noView, then read access will not be granted.

WriteViewName

```
Synopsis: { noView, V1Mib, allOfMib }
Default: noView
```

This parameter identifies the MIB tree(s) to which this entry authorizes write access. If the value is noView, then write access will not be granted.

NotifyViewName

```
Synopsis: { noView, V1Mib, allOfMib }
Default: noView
```

This parameter identifies the MIB tree(s) to which this entry authorizes access for notifications. If the value is noView, then access for notifications will not be granted.

1.12. RADIUS

RADIUS (Remote Authentication Dial In User Service) is used to provide centralized authentication and authorization for network access. ROS assigns a privilege level of Admin, Operator or Guest to a user who presents a valid user name and password. The number of users who can access the ROS server is ordinarily dependent on the number of user records which can be configured on the server itself. ROS can also, however, be configured to pass along the credentials provided by the user to be remotely authenticated by a RADIUS server. In this way, a single RADIUS server can centrally store user data and provide authentication and authorization service to multiple ROS servers needing to authenticate connection attempts.

1.12.1. RADIUS overview

RADIUS (described in RFC 2865 [http://tools.ietf.org/html/rfc2865]) is a UDP-based protocol is used for carrying authentication, authorization, and configuration information between a Network Access Server which desires to authenticate its links and a shared Authentication Server. RADIUS is also used also widely utilized in conjunction with 802.1x for port security using EAP (See Appendix A).

A RADIUS server can act as a proxy client to other RADIUS servers or other kinds of authentication servers.

Unlike TACACS+, authorization and authentication functionality is supported by RADIUS in the same packet frame. TACACS+ actually separates authentication from authorization into separate packets.

On receiving an authentication-authorization request from a client in an “Access-Request” packet, the RADIUS server checks the conditions configured for received username-password combination in the user database. If all the conditions are met, the list of configuration values for the user is placed into an “Access-Accept” packet. These values include the type of service (e.g. SLIP, PPP, Login User) and all the necessary values to deliver the desired service.

1.12.2. User Login Authentication and Authorization

A RADIUS server can be used to authenticate and authorize access to the device’s services, such as HMI via Serial Console, Telnet, SSH, RSH, Web Server (see Password Configuration). ROS implements a RADIUS client which uses the Password Authentication Protocol (PAP) to verify access. Attributes sent to a RADIUS server are:

- user name
- user password
- service type: Login
- vendor specific, currently defined as the following:

vendor ID: RuggedCom Inc. enterprise number (15004) assigned by the Internet Assigned Numbers Authority (IANA)

string, sub-attribute containing specific values:

subtype: 1 (vendor's name subtype)

length: 11 (total length of sub-attribute of subtype 1)

ASCII string "RuggedCom"

Two RADIUS servers (Primary and Secondary) are configurable per device. If the Primary Server is not reachable, the device will automatically fall back to the Secondary server to complete the authorization process.

The vendor specific attribute is used to determine the access level from the server, which may be configured at the RADIUS server with the following information:

- Vendor ID: RuggedCom Inc. enterprise number (15004) assigned by Internet Assigned Numbers Authority (IANA)
- Sub-attribute Format: String
- Vendor Assigned Sub-Attribute Number: 2
- Attribute value – any one of: admin, operator, guest

Note

If no access level is received in the response packet from the server then no access will be granted to the user

---

An Example of a RuggedCom Dictionary for a FreeRADIUS server:

VENDOR	RuggedCom 15004
BEGIN-VENDOR	RuggedCom
ATTRIBUTE	RuggedCom-Privilege-level 2 string
END-VENDOR	RuggedCom

Sample entry for user "admin" Adding Users:

admin Auth-Type := Local, User-Password == "admin"

RuggedCom-Privilege-level = "admin

1.12.3. 802.1X Authentication

A RADIUS server may also be used to authenticate access on ports with 802.1X security support. Attributes sent to the RADIUS server in a RADIUS Request are:

- user name, derived from client's EAP identity response
- NAS IP address
- service type: framed
- framed MTU:1500 (maximum size of EAP frame, which is the size of an Ethernet frame)

- EAP message
- vendor specific attribute, as described above

RADIUS messages are sent as UDP messages. The switch and the RADIUS server must use the same authentication and encryption key.

1.12.4. Radius Server Configuration

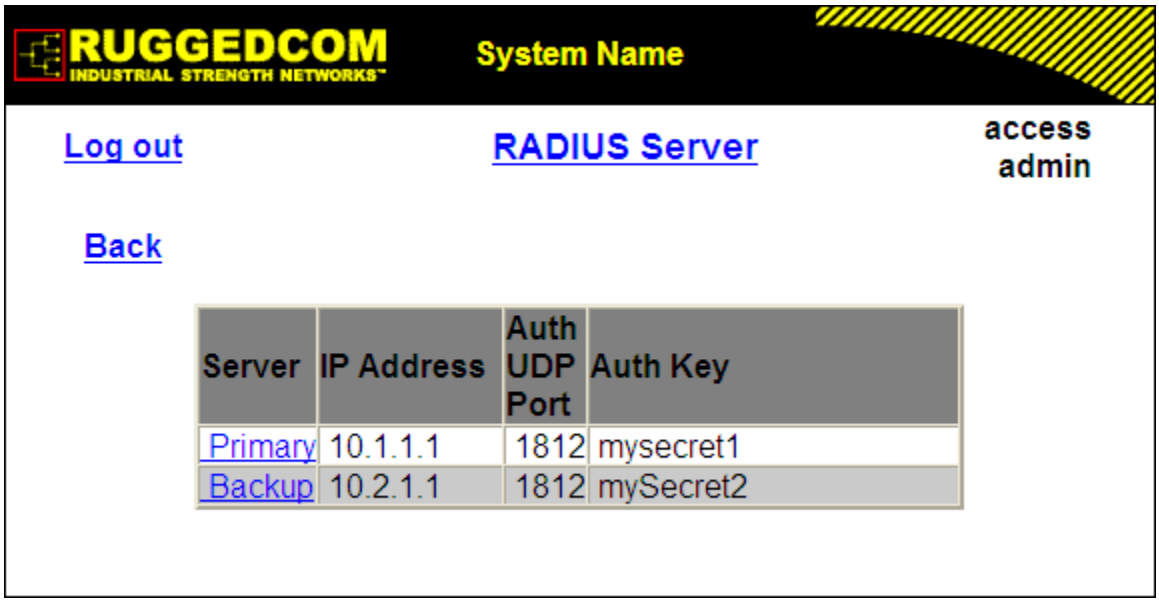

Figure 1.25. RADIUS Server Summary

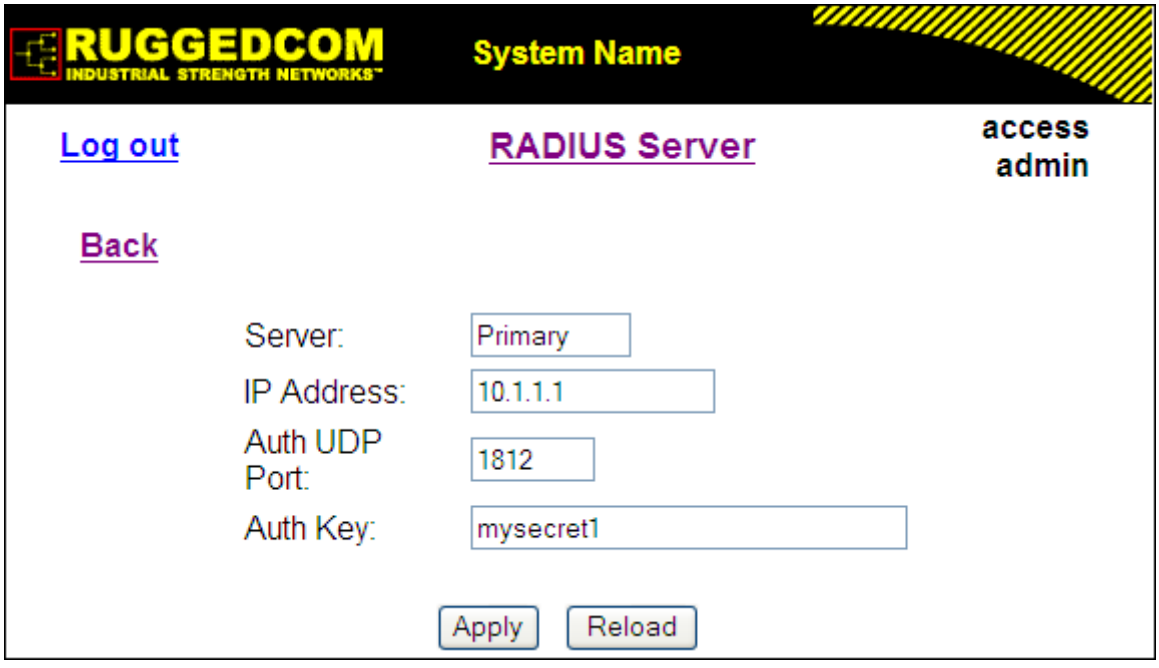

Figure 1.26. RADIUS Server Form

Server

```
Synopsis: Any 8 characters
```

```
Default: Primary
```

This field tells whether this configuration is for a primary or a backup server

IP Address

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default:
```

The RADIUS server IP Address.

Auth UDP Port

```
Synopsis: 1 to 65535
Default: 1812
```

The authentication UDP Port on the RADIUS server.

Auth Key

```
Synopsis: 31 character ASCII string
Default: None
```

The authentication key shared with the RADIUS server. It is used to encrypt any passwords that are sent between the switch and the RADIUS server.

Confirm Auth Key

```
Synopsis: 31 character ASCII string
Default: None
```

Confirm input of the above authentication key.

1.13. TACACS+

TACACS+ (Terminal Access Controller Access-Control System Plus) is a TCP-based access control protocol that provides authentication, authorization and accounting services to routers, network access servers and other networked computing devices via one or more centralized servers. It is based on, but is not compatible with, the older TACACS protocol. TACACS+ has generally replaced its predecessor in more recently built or updated networks, although TACACS and XTACACS are still used on many older networks. Note that RuggedCom's TACACS+ client implementation always has encryption enabled.

1.13.1. User Login Authentication and Authorization

A TACACS+ server can be used to authenticate and authorize access to the device's services, such as HMI via Serial Console, Telnet, SSH, RSH, Web Server (see Password Configuration). User name and Password are sent to the configured TACACS+ Server.

Two TACACS+ servers (Primary and Secondary) are configurable per device. If the primary server is not reachable, the device will automatically fall back to the secondary server to complete the authorization process.

- The TACACS+ standard priv_lvl attribute will be used to grant access to the device:

  priv_lvl=15 represents an access level of "admin"

  1 < priv_lvl < 15 represents an access level of "operator" (i.e. any value from 2 to 14)

  priv_lvl=1 represents an access level of "guest"

Note

If an access level is not received in the response packet from the server, then access will not be granted to the user.

1.13.2. TACACS+ Server Configuration

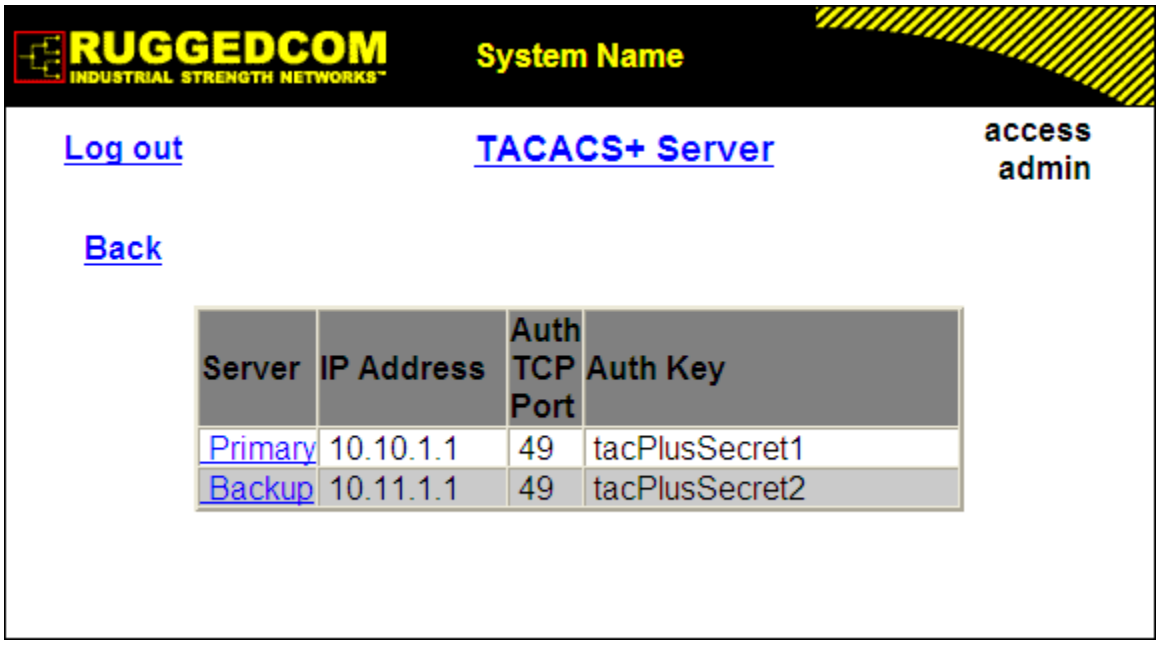

Figure 1.27. TACACS+ Server Summary

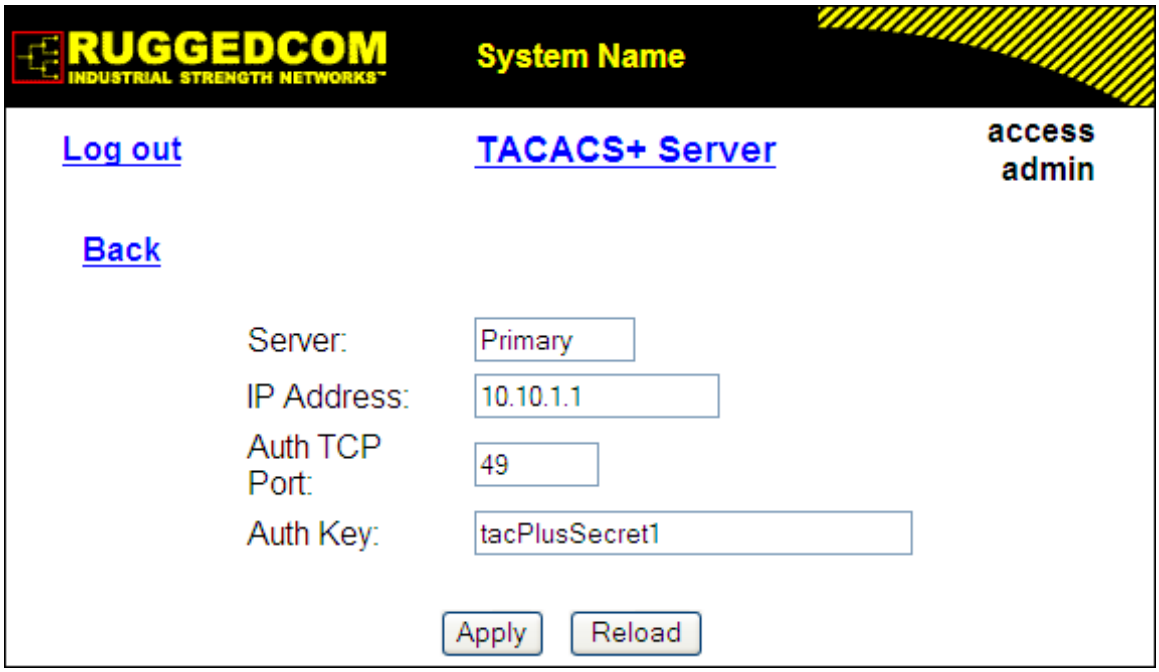

Figure 1.28. TACACS+ Server Form

Server

```
Synopsis: Any 8 characters
Default: Primary
```

This field indicates whether this configuration is for a primary or a backup server.

IP Address

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default:
```

The TACACS+ server IP Address.

Auth TCP Port

```
Synopsis: 1 to 65535
Default: 49
```

The authentication TCP Port on the TACACS+ server.

Auth Key

```
Synopsis: 31 character ASCII string
Default:
```

The authentication key shared with the TACACS+ server. It is used to encrypt any passwords that are sent from the switch to the TACACS+ server.

Confirm Auth Key

```
Synopsis: 31 character ASCII string
Default: None
```

Confirm input of the above authentication key.

1.14. DHCP Relay Agent

A DHCP Relay Agent is a device that forwards DHCP packets between clients and servers when they are not on the same physical LAN segment or IP subnet. The feature is enabled if the DHCP server IP address and a set of access ports are configured.

DHCP Option 82 provides a mechanism for assigning an IP Address based on the location of the client device in the network. Information about the client's location can be sent along with the DHCP request to the server. The DHCP server makes a decision about an IP Address to be assigned, based on this information.

DHCP Relay Agent takes the broadcast DHCP requests from clients received on the configured access port and inserts the relay agent information option (Option 82) into the packet. Option 82 contains the VLAN ID (2 bytes) and the port number of the access port (2 bytes - the circuit ID sub-option) and the switch's MAC address (the remote ID sub-option). This information uniquely defines the access port's position in the network.

The DHCP Server supporting DHCP option 82 sends a unicast reply and echoes Option 82. The DHCP Relay Agent removes the Option 82 field and broadcasts the packet to the port from which the original request was received.

These parameters provide the ability to configure the switch to act as a relay agent for DHCP Option 82.

The DHCP Relay Agent is communicating to the server on a management interface. The agent's IP address is the address configured for the management interface.

Figure 1.29. DHCP Relay Agent Form

DHCP Server Address

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default:
```

This parameter specifies the IP address of the DHCP server to which DHCP queries will be forwarded from this relay agent.

DHCP Client Ports

```
Synopsis: Any combination of numbers valid for this parameter
Default: None
```

This parameter specifies ports where DHCP clients are connected.

Examples:

- **All** - all ports of the switch can have DHCP clients connected.
- **2,4-6,8** - ports 2,4,5,6 and 8 can have DHCP clients connected

1.15. Syslog

The syslog provides users with the ability to configure local and remote syslog connections. The remote syslog protocol, defined in RFC 3164, is a UDP/IP-based transport that enables a device to send event notification messages across IP networks to event message collectors, also known as syslog servers. The protocol is simply designed to transport these event messages from the generating device to the collector.

The syslog client resides in ROS and supports up to 5 collectors (syslog servers). ROS Remote Syslog provides the ability to configure:

- IP address(es) of collector(s).
- Source UDP port.
- Destination UDP port per collector.
- Syslog source facility ID per collector (same value for all ROS modules).

- Filtering severity level per collector (in case different collectors are interested in syslog reports with different severity levels).

1.15.1. Configuring Local Syslog

The local syslog configuration enables users to control what level of syslog information will be logged. Only messages of a severity level equal to or greater than the configured severity level are written to the syslog.txt file in the unit.

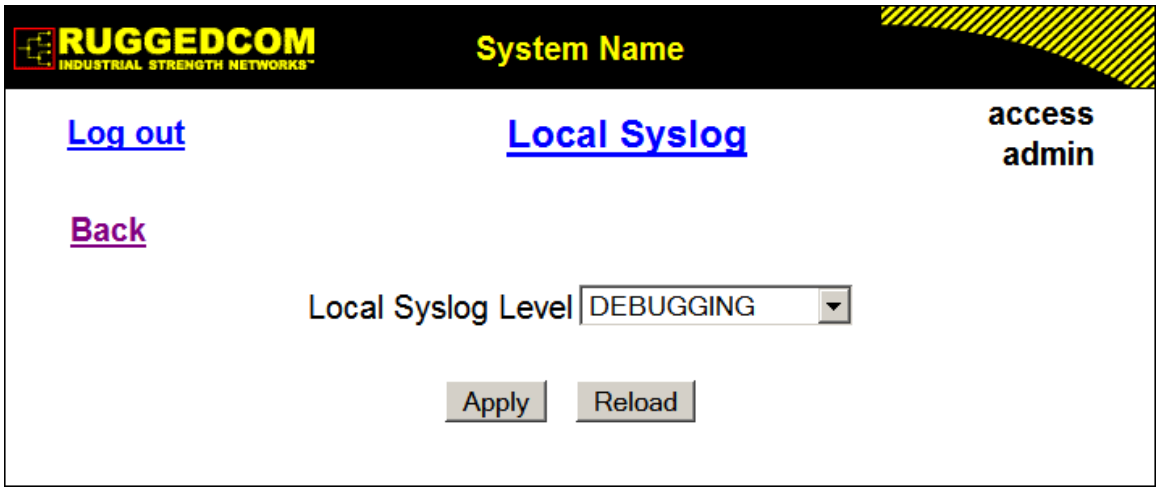

Figure 1.30. Local Syslog Form

Local Syslog Level

```
Synopsis: { EMERGENCY, ALERT, CRITICAL, ERROR, WARNING, NOTICE,
INFORMATIONAL, DEBUGGING }
Default: DEBUGGING
```

Syslog severity level - {EMERGENCY, ALERT, CRITICAL, ERROR, WARNING, NOTICE, INFORMATIONAL, DEBUGGING}.

1.15.2. Configuring Remote Syslog Client

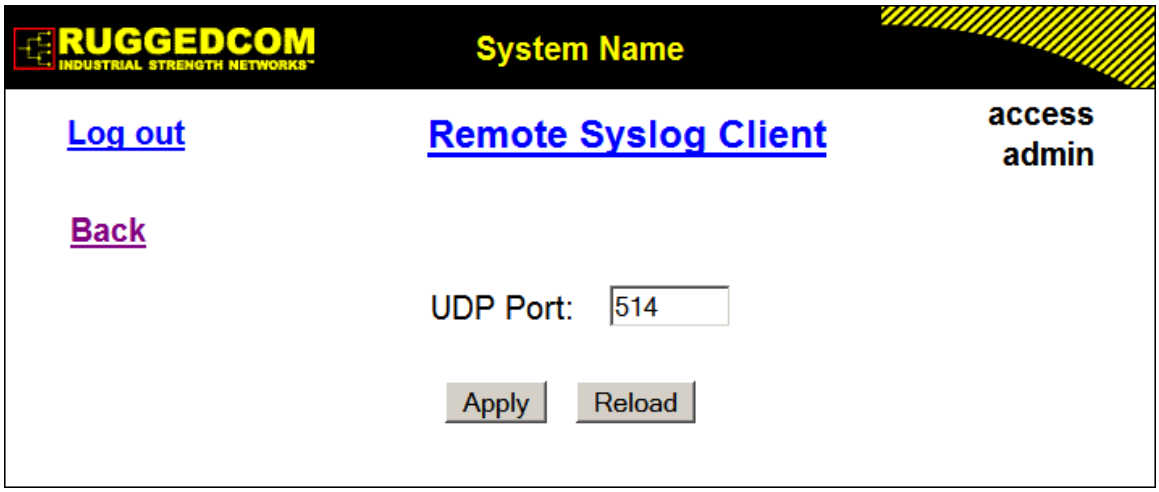

Figure 1.31. Remote Syslog Client Form

UDP Port

```
Synopsis: 1025 to 65535 or { 514 }
Default: 514
```

The local UDP port through which the client sends information to the server(s).

1.15.3. Configuring the Remote Syslog Server

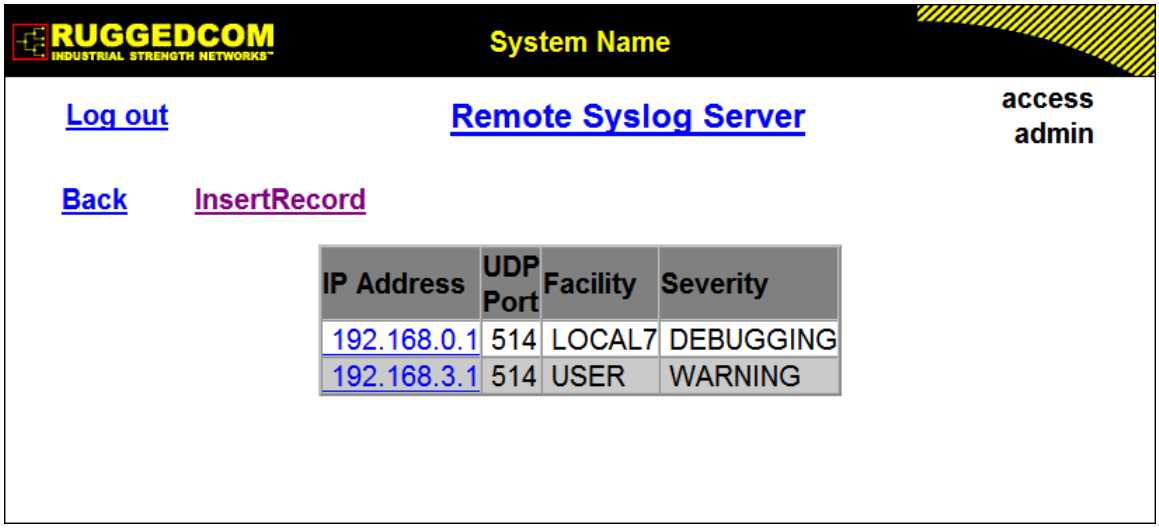

Figure 1.32. Remote Syslog Server Table

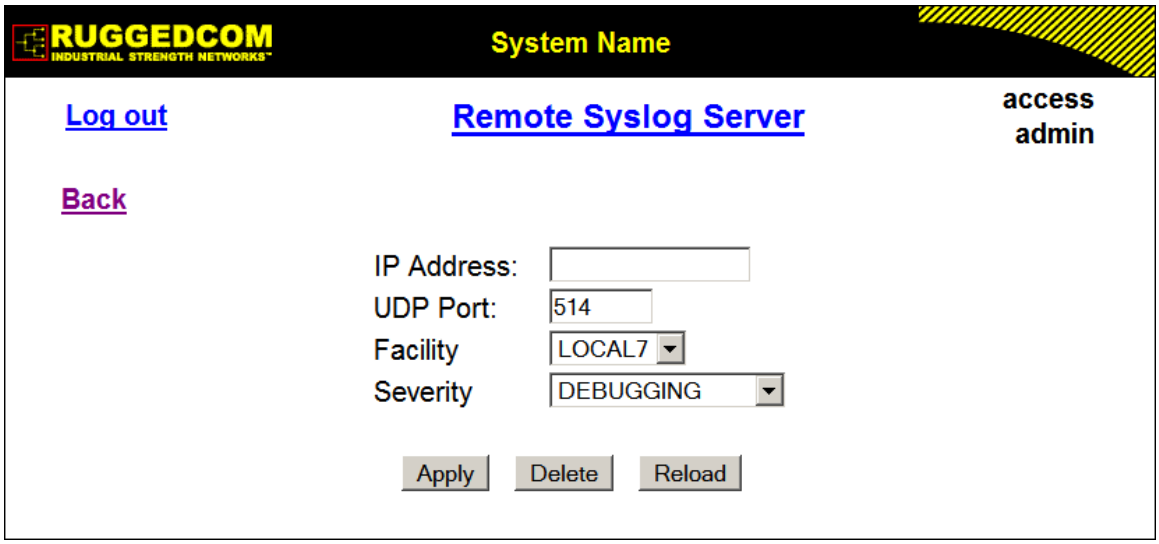

Figure 1.33. Remote Syslog Server Form

IP Address

```
Synopsis: ###.###.###.### where ### ranges from 0 to 255
Default:
```

Syslog server IP Address.

UDP Port

```
Synopsis: 1025 to 65535 or { 514 }
```

`Default: 514`

The UDP port number on which the remote server listens.

Facility

`Synopsis: { USER, LOCAL0, LOCAL1, LOCAL2, LOCAL3, LOCAL4, LOCAL5, LOCAL6, LOCAL7 }`

`Default: LOCAL7`

Syslog facility name - { USER, LOCAL0, LOCAL1, LOCAL2, LOCAL3, LOCAL4, LOCAL5, LOCAL6, LOCAL7 }.

Severity

`Synopsis: { EMERGENCY, ALERT, CRITICAL, ERROR, WARNING, NOTICE, INFORMATIONAL, DEBUGGING }`

`Default: DEBUGGING`

Syslog severity level - {EMERGENCY, ALERT, CRITICAL, ERROR, WARNING, NOTICE, INFORMATIONAL, DEBUGGING}.

1.16. Troubleshooting

Problem One

I have configured the IP address and a gateway. I am pinging the switch but it is not responding. I am sure the switch is receiving the ping because its port LEDs are flashing and the statistics menu shows the pings. What is going on?

Is the switch being pinged through a router? If so, the switch gateway address must be configured. The following figure illustrates the problem.

Figure 1.34. Using A Router As A Gateway

The router is configured with the appropriate IP subnets and will forward the ping from the workstation to the switch. When the switch responds, however, it will not know which of its interfaces to use in order to reach the workstation and will drop the response. Programming a gateway of 10.0.0.1 will cause the switch to forward unresolvable frames to the router.

This problem will also occur if the gateway address is not configured and the switch tries to raise an SNMP trap to a host that is not on the local subnet.

2. Ethernet Ports

ROS® Ethernet port control provides the following features:

- Configuring port physical parameters.
- Configuring link alarms/traps for the port.
- Configuring port rate limiting.
- Using Port Mirroring.
- Cable Diagnostics.
- Viewing port status.
- Resetting all or some ports.
- Using Link-Fault-Indication (LFI).

2.1. Controller Protection Through Link-Fault-Indication (LFI)

Modern industrial controllers often feature backup Ethernet ports used in the event of a link failure. When these interfaces are supported by media (such as fiber) that employ separate transmit and receive paths, the interface can be vulnerable to failures that occur in only one of the two paths.

Refer to the following figure. While the link between switch A and the controller functions normally, the controller holds the backup link down. Switch B learns that it must forward frames towards switch A in order to reach the controller.

Unfortunately, if the transmission path from the controller to switch A fails, switch A will still generate link signals to the controller. The controller will still detect link to switch A and will not fail over to the backup port.

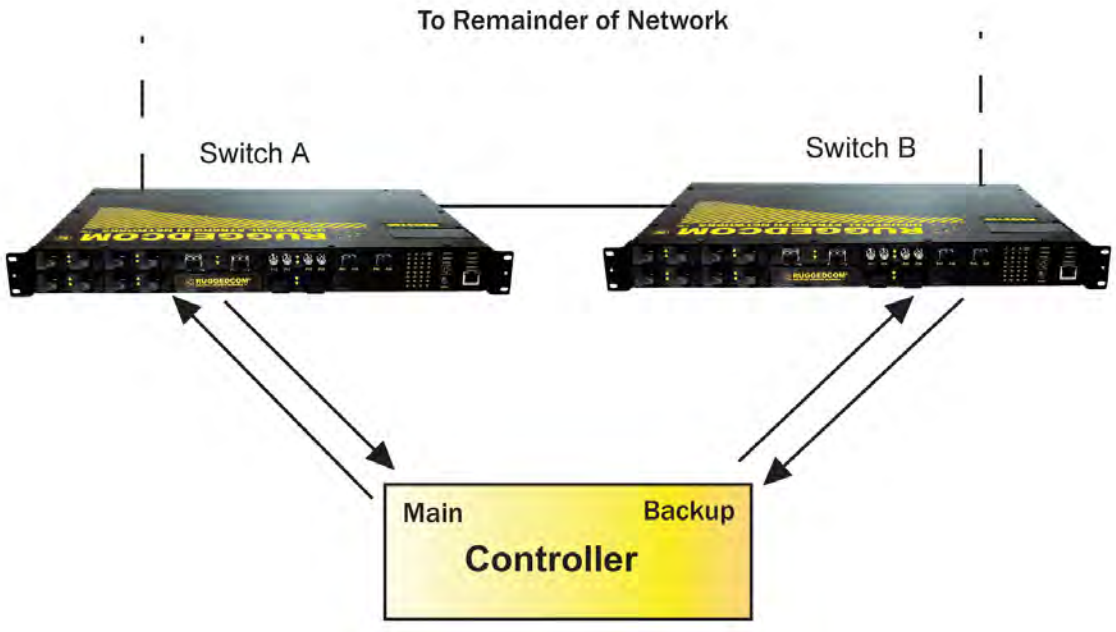

Figure 2.1. Controller Protection Through LFI

To overcome this problem, there should be a way of notifying the link partner in case a link integrity signal stopped being received from it. Such a way natively exists in some link media but not in others:

- **Auto-Negotiating links (100Base-TX,1000Base-T,1000Base-X)** - auto-negotiation built-in feature (a special flag called Remote Fault Indication is set in the transmitted auto-negotiation signal)
- **100Base-FX links** - Far–End-Fault-Indication (FEFI) is a standard feature defined by the IEEE 802.3 standard for this link type. The feature includes:
  - Transmitting FEFI - transmitting modified link integrity signal in case a link failure is detected, i.e. no link signal is received from the link partner.
  - Detecting FEFI - indicating link loss in case FEFI signal is received from the link partner.
- **10Base-FL links** - no standard support

As one can see from the above, 10Base-FL links have no native link partner notification mechanism. Also, FEFI support in 100Base-FX links is optional according to the IEEE 802.3 standard, which means that some link partners may not support it.

RuggedCom offers an advanced Link-Fault-Indication (LFI) feature for the links where no native link partner notification mechanism is available. With the LFI enabled, the device bases generation of a link integrity signal upon its reception of a link signal. In the diagram above, if switch A fails to receive a link signal from the controller, it will stop generating a link signal. The controller will detect the link failure and switch to the backup port.

The switch can also be configured to flush the MAC address table for the controller port (see MAC Address Tables section). Frames destined for the controller will be flooded to switch B where they will be forwarded to the controller (after the controller transmits its first frame).

Note

If both link partners are capable of the LFI, it MUST NOT be enabled on both sides of the link. If it is enabled on both sides, the link will never be established because each side will permanently wait for its partner to transmit a link signal.

2.2. Ethernet Ports Configuration and Status

The Ethernet Ports menu is accessible from the main menu.

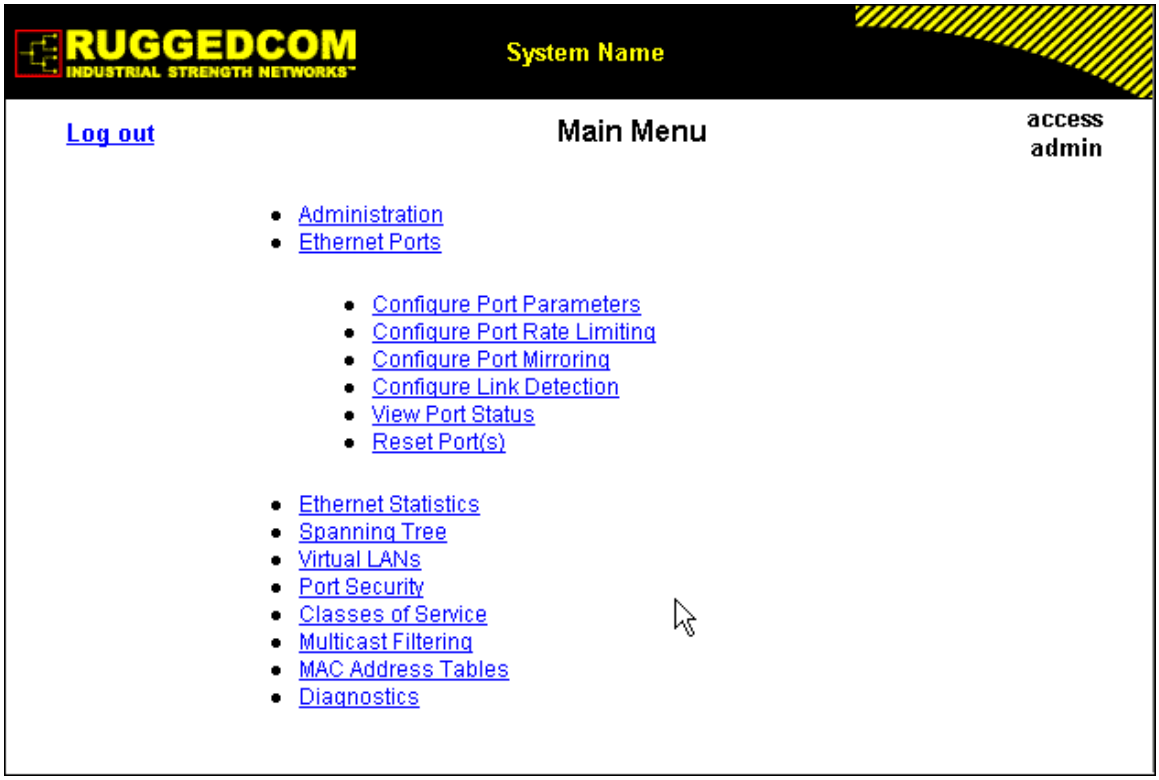

Figure 2.2. Ethernet Ports Menu

2.2.1. Port Parameters

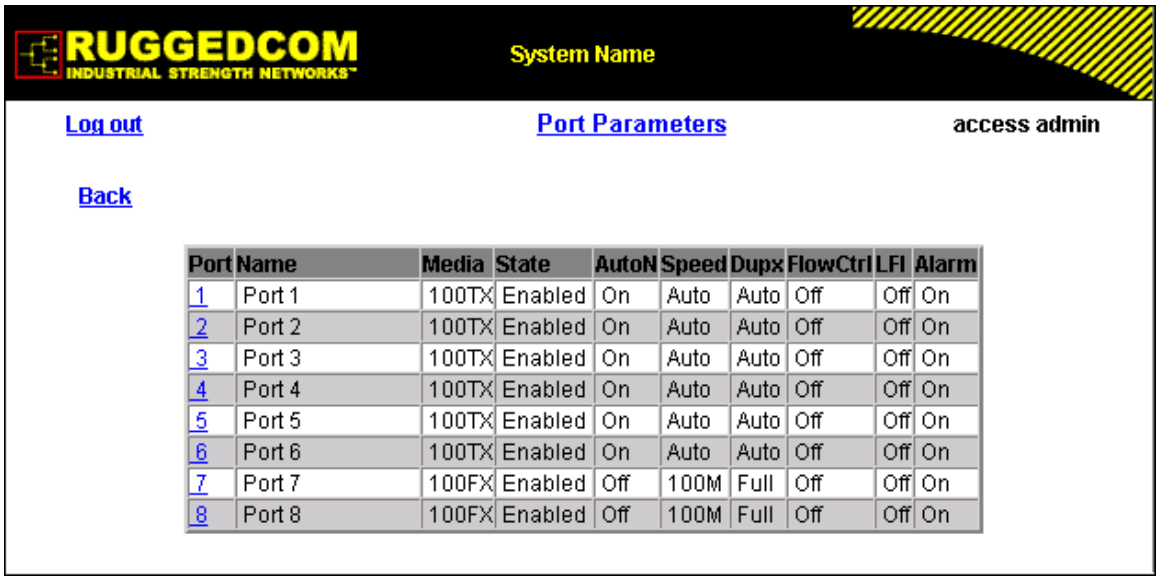

Port	Name	Media	State	AutoN	Speed	Dupx	FlowCtrl	LFI	Alarm
1	Port 1	100TX	Enabled	On	Auto	Auto	Off	Off	On
2	Port 2	100TX	Enabled	On	Auto	Auto	Off	Off	On
3	Port 3	100TX	Enabled	On	Auto	Auto	Off	Off	On
4	Port 4	100TX	Enabled	On	Auto	Auto	Off	Off	On
5	Port 5	100TX	Enabled	On	Auto	Auto	Off	Off	On
6	Port 6	100TX	Enabled	On	Auto	Auto	Off	Off	On
7	Port 7	100FX	Enabled	Off	100M	Full	Off	Off	On
8	Port 8	100FX	Enabled	Off	100M	Full	Off	Off	On

Figure 2.3. Port Parameters Table

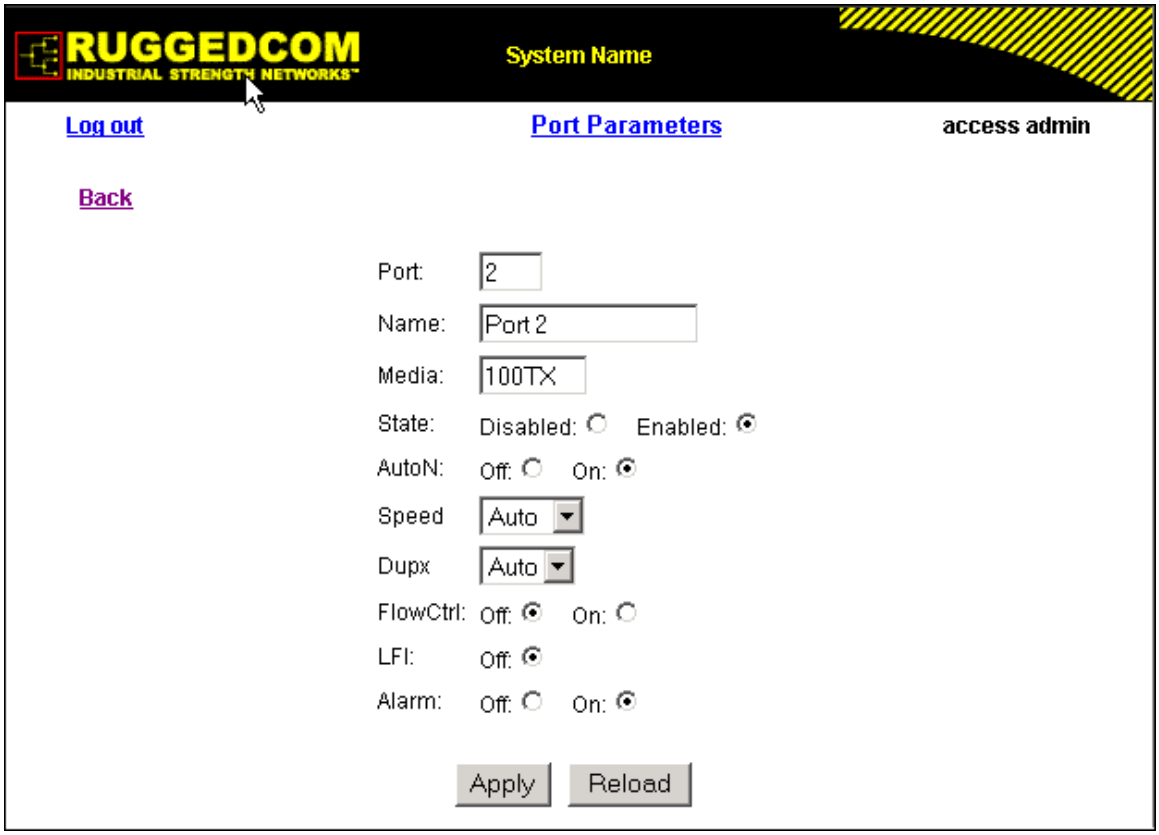

Figure 2.4. Port Parameters Form

Port

```
Synopsis: 1 to maximum port number
Default: 0
```

The port number as seen on the front plate silkscreen of the switch.

Name

```
Synopsis: Any 15 characters
Default: Not installed
```

A descriptive name that may be used to identify the device connected to that port.

Media

```
Synopsis: { 100TX, 10FL, 100FX, 1000X, 1000T }
```

The type of the port's media.

State

```
Synopsis: { Disabled, Enabled }
Default: Enabled
```

Disabling a port will prevent all frames from being sent and received on that port. Also, when disabled link integrity pulses are not sent so that the link/activity LED will never be lit. You may want to disable a port for troubleshooting or to secure it from unauthorized connections.

AutoN

```
Synopsis: { Off, On }
Default: On
```

Enable or disable IEEE 802.3 auto-negotiation. Enabling auto-negotiation results in speed and duplex mode being negotiated upon link detection; both end devices must be auto-negotiation compliant for the best possible results. 10Mbps and 100Mbps fiber optic media do not support auto-negotiation so these media must be explicitly configured to either half or full-duplex mode. Full-duplex operation requires both ends to be configured as such or else severe frame loss will occur during heavy network traffic.

Speed

```
Synopsis: { Auto, 10M, 100M, 1G }
Default: Auto
```

Speed (in Megabit-per-second or Gigabit-per-second). If auto-negotiation is enabled, this is the speed capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is set to this speed.

AUTO means advertise all supported speed modes.

Dupx

```
Synopsis: { Auto, Half, Full }
Default: Auto
```

Duplex mode. If auto-negotiation is enabled, this is the duplex capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is set to this duplex mode.

AUTO means advertise all supported duplex modes.

Flow Control

```
Synopsis: { Off, On }
Default: Off
```

Flow Control is useful for preventing frame loss during times of severe network traffic. Examples of this include multiple source ports sending to a single destination port or a higher-speed port bursting to a lower-speed port.

When the port is in half-duplex mode, it is accomplished using 'backpressure' where the switch simulates collisions causing the sending device to retry transmissions according to the Ethernet back-off algorithm. When the port is full-duplex it is accomplished using PAUSE frames which cause the sending device to stop transmitting for a certain period of time.

LFI

```
Synopsis: { Off, On }
Default: Off
```

Enabling Link-Fault-Indication (LFI) inhibits transmission of the link integrity signal when the receiving link has failed. This enables the device at far end to detect link failure under all circumstances.

Note

This feature must not be enabled at both ends of a link.

Link Alarms

```
Synopsis: { Off, On }
Default: On
```

Disabling link state alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that port.

Note

If one end of the link is fixed to a specific speed and duplex type and the peer auto-negotiates, there is a strong possibility that the link will either fail to raise, or raise with the wrong settings on the auto-negotiating side. The auto-negotiating peer will fall back to half-duplex operation, even when the fixed side is full duplex. Full-duplex operation requires that both ends are configured as such or else severe frame loss will occur during heavy network traffic. At lower traffic volumes the link may display few if any errors As the traffic volume rises the fixed negotiation side will begin to experience dropped packets while the auto-negotiating side will experience excessive collisions. Ultimately, as traffic load approaches 100% the link will become entirely unusable. These problems can be avoided by always configuring ports to the appropriate fixed values.

2.2.2. Port Rate Limiting

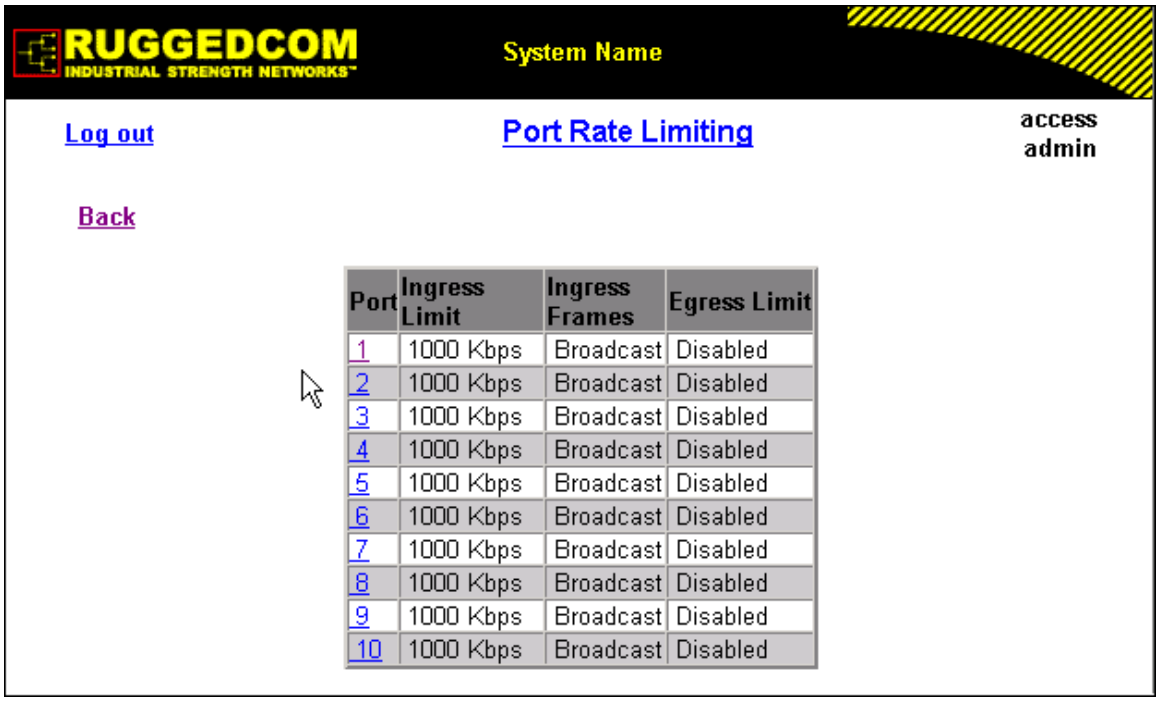

Port	Ingress Limit	Ingress Frames	Egress Limit
1	1000 Kbps	Broadcast	Disabled
2	1000 Kbps	Broadcast	Disabled
3	1000 Kbps	Broadcast	Disabled
4	1000 Kbps	Broadcast	Disabled
5	1000 Kbps	Broadcast	Disabled
6	1000 Kbps	Broadcast	Disabled
7	1000 Kbps	Broadcast	Disabled
8	1000 Kbps	Broadcast	Disabled
9	1000 Kbps	Broadcast	Disabled
10	1000 Kbps	Broadcast	Disabled

Figure 2.5. Port Rate Limiting Table

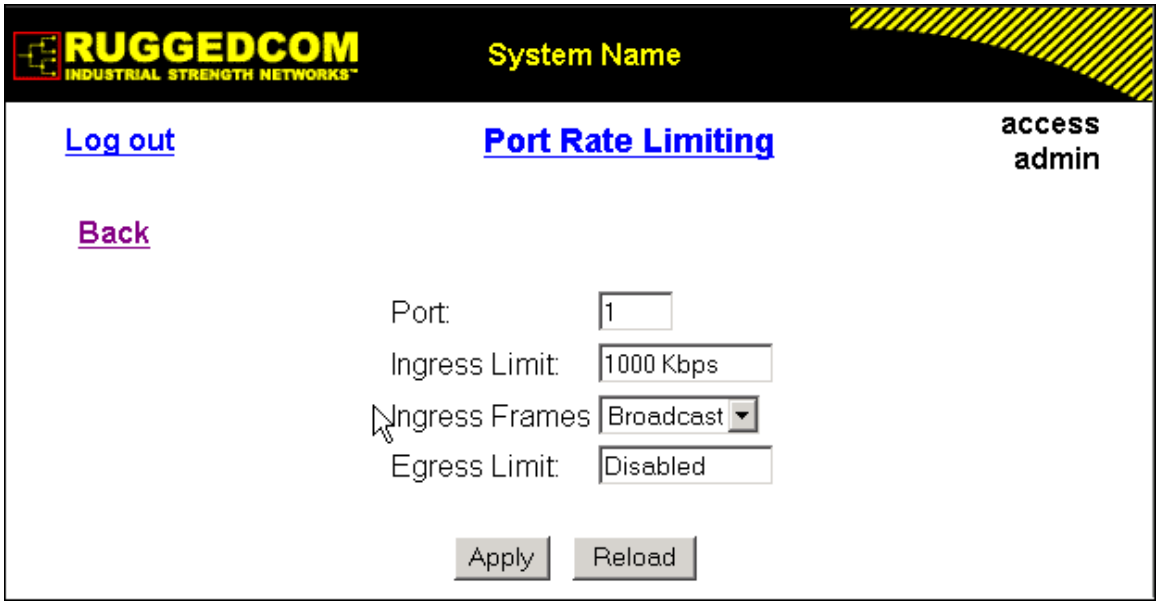

Figure 2.6. Port Rate Limiting Form

Port

```
Synopsis: 1 to maximum port number
Default: 1
```

The port number as seen on the front plate silkscreen of the switch.

Ingress Limit

```
Synopsis: 62 to 256000 Kbps or { Disabled }
Default: 1000 Kbps
```

The rate at which received frames (of the type described by the ingress frames parameter) will start to be discarded by the switch.

Ingress Frames

```
Synopsis: { Broadcast, Multicast, All }
Default: Broadcast
```

This parameter specifies the types of frames to rate-limit on this port. It applies only to received frames:

BROADCAST - only broadcast frames will be limited

MULTICAST - all multicast frames (including broadcast) will be limited

ALL - all frames (both multicast and unicast) will be limited

Egress Limit

```
Synopsis: 62 to 256000 Kbps or { Disabled }
Default: Disabled
```

The maximum rate at which the switch will transmit (multicast, broadcast and unicast) frames on this port. The switch will discard frames in order to meet this rate if required.

2.2.3. Port Mirroring

Port mirroring is a troubleshooting tool that copies, or mirrors, all traffic received or transmitted on a designated port to another mirror port. If a protocol analyzer were attached to the target port, the traffic stream of valid frames on any source port is made available for analysis.

Select a target port that has a higher speed than the source port. Mirroring a 100 Mbps port onto a 10 Mbps port may result in an improperly mirrored stream.

Frames will be dropped if the full-duplex rate of frames on the source port exceeds the transmission speed of the target port. Since both transmitted and received frames on the source port are mirrored to the target port, frames will be discarded if the sum traffic exceeds the target port's transmission rate. This problem reaches its extreme in the case where traffic on a 100 Mbps full-duplex port is mirrored onto a 10 Mbps half-duplex port.

Note

Invalid frames received on the source port will not be mirrored. These include CRC errors, oversize and undersize packets, fragments, jabbers, collisions, late collisions and dropped events).

2.2.3.1. Port Mirroring Limitations

- Traffic will be mirrored onto the target port only if the target port is a member of the same VLANs as the source port.
- The target port may sometimes incorrectly show the VLAN tagged/untagged format of the mirrored frames.
- Network management frames (such as RSTP, GVRP etc.) may not be mirrored.
- Switch management frames generated by the switch (such as Telnet, HTTP, SNMP etc.) may not be mirrored.

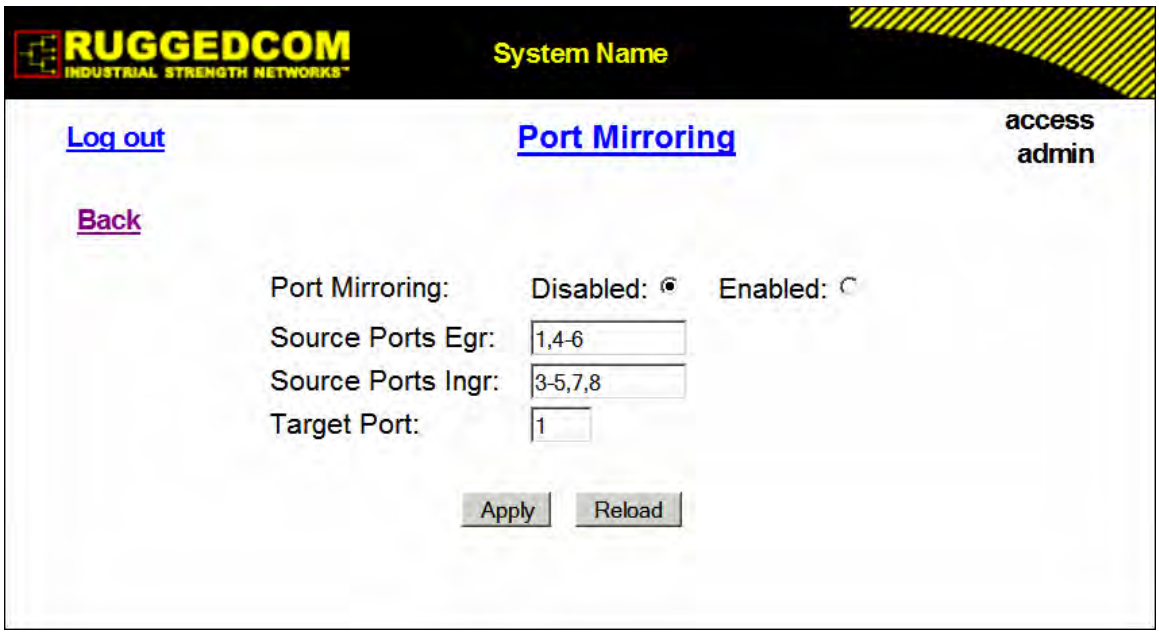

Figure 2.7. Port Mirroring Form

Port Mirroring

```
Synopsis: { Disabled, Enabled }
Default: Disabled
```

Enabling port mirroring causes all frames received and transmitted by the source port(s) to be transmitted out of the target port.

Source Ports Egr

```
Synopsis: Any combination of numbers valid for this parameter
Default: None
```

Ethernet ports whose egress traffic is to be mirrored to the target port.

Source Ports Ingr

```
Synopsis: Any combination of numbers valid for this parameter
Default: None
```

Ethernet ports whose ingress traffic is to be mirrored to the target port.

Target Port

```
Synopsis: 1 to maximum port number
Default: 1
```

The port to which selected traffic is mirrored. A monitoring device should be connected to the target port.

2.2.4. Cable Diagnostics

ROS is able to perform cable diagnostics per Ethernet port and to view the results.

Warning

When cable diagnostics are performed on a port, any established network link on the port will be dropped and normal network traffic will not be able to pass through either the Port Under Test or the Partner Port. Please be aware of the potential network interruption that could be triggered by running cable diagnostics. After the cable diagnostics finish, the original network port settings for both the Port Under Test and the Partner Port are restored along with any established link.

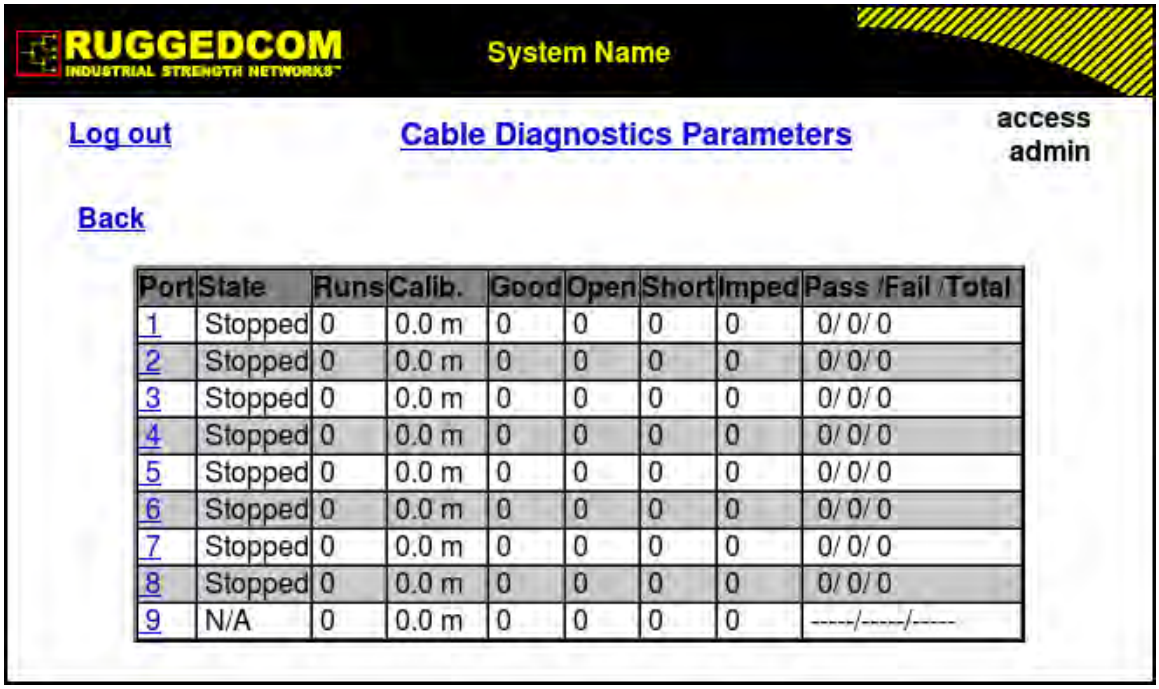

Port	State	Runs	Calib.	Good	Open	Short	Imped	Pass /Fail /Total
1	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
2	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
3	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
4	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
5	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
6	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
7	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
8	Stopped	0	0.0 m	0	0	0	0	0/ 0/ 0
9	N/A	0	0.0 m	0	0	0	0	-----/-----/-----

Figure 2.8. Cable Diagnostics Table

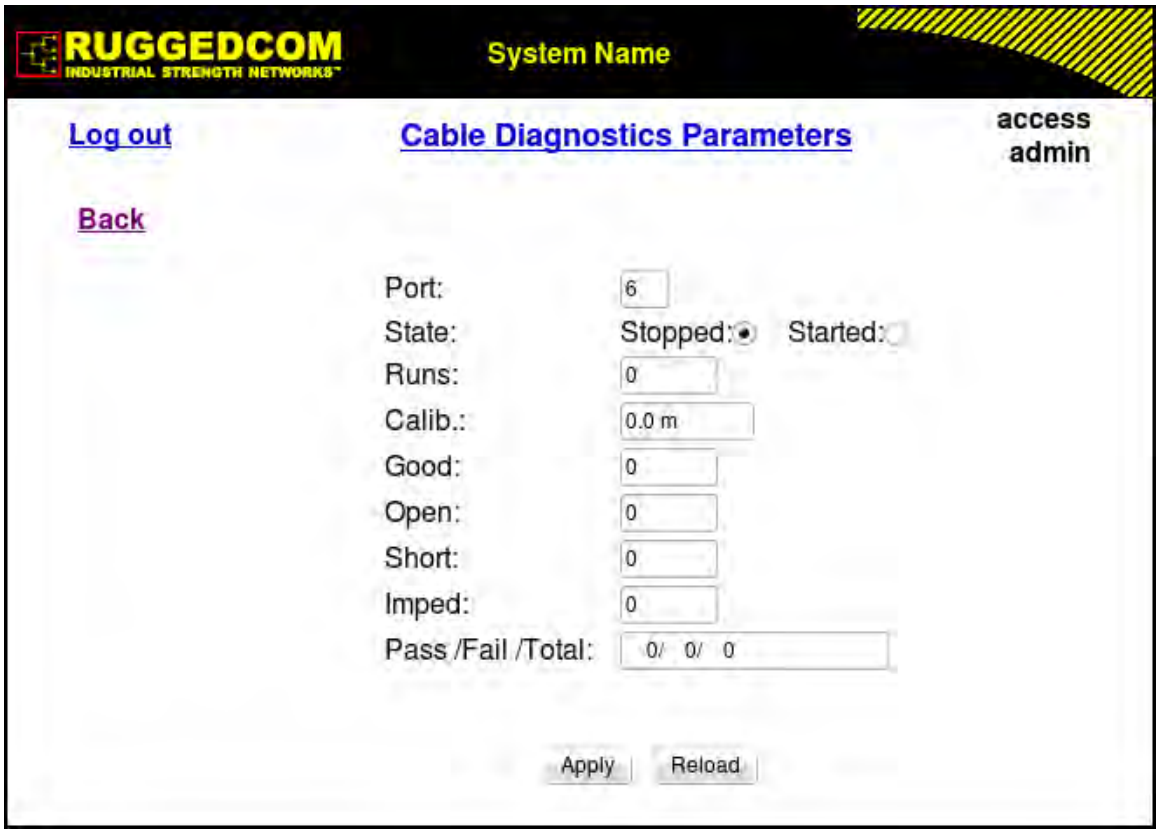

Figure 2.9. Cable Diagnostics Parameters Form

The Cable Diagnostics Table screen, pictured above, lists the current value of the following parameters for all Ethernet ports. Clicking on a port number in the table brings up the Cable

Diagnostics Form for the corresponding port. This form can be used to set certain of the cable diagnostic parameters for the port, as indicated below:

Port

`Synopsis: 1 to X`

The port number as seen on the front plate silkscreen of the switch.

State

`Started, Stopped or N/A`

Start or stop cable diagnostics on the selected port. If a port does not support cable diagnostics, State will be reported as N/A.

Runs

`Synopsis: 0 to 65535`

The total number of times that cable diagnostics are to be performed on the selected port. If set to 0, cable diagnostics will be performed until diagnostics are stopped explicitly.

Calib.

`Synopsis: -100.0 m to 100.0 m`

The calibration value can be used to adjust the estimated distance to the fault. Refer to Calibrating Estimated Distance To Fault for details on setting this parameter.

Good

`Synopsis: 0 to 65535`

The number of times that GOOD TERMINATION (no fault) has been detected on the cable pairs of the selected port.

Open:

`Synopsis: 0 to 65535`

The number of times that OPEN has been detected on the cable pairs of the selected port.

Short

`Synopsis: 0 to 65535`

The number of times that SHORT has been detected on the cable pairs of the selected port.

Imped

`Synopsis: 0 to 65535`

The number of times that IMPEDANCE MISMATCH has been detected on the cable pairs of the selected port.

Pass/Fail/Total:

`Synopsis: 0 to 65535 / 0 to 65535 / 0 to 65535`

This field summarizes the results of the cable diagnostics performed so far:

- Pass - the number of times that cable diagnostics were completed successfully on the selected port.
- Fail - the number of times that cable diagnostics failed on the selected port.
- Total - the total number of times that cable diagnostics have been attempted on the selected port.

2.2.4.1. Running Cable Diagnostics

To start cable diagnostics on a port:

1. Connect a Category 5 or better quality cable to the port under test (PUT).
2. Connect the other end of the cable to a similar network port. For example, connect 100BASE-T port to a 100BASE-T port, 1000BASE-T port to a 1000BASE-T port.
3. Configure the PUT's "Runs" count.
4. Configure the PUT's cable diagnostics State to "Started".

To stop cable diagnostics on a port:

1. Configure the PUT's cable diagnostics state to "Stopped". Diagnostics may be stopped at any point. If a stop is issued in the middle of a diagnostics run, it will nevertheless run to completion and the results will be updated.

Note

Both the port under test (PUT) or partner port (PT) can be configured to be either in Enabled mode with auto-negotiation or in Disabled mode. Other modes may interfere with the cable diagnostics procedure and are not recommended.

2.2.4.2. Interpreting Cable Diagnostics Results

Four different conditions are reported for the state of a cable under examination:

- Good - No fault is detected on the tested cable.
- Open - Opened cable pair(s) is/are detected on the tested cable.
- Short - Short cable pair(s) is/are detected on the tested cable.
- Imped - Impedance Mismatch is detected on the tested cable.

The corresponding counts for each of these status conditions indicates the number of occurrences of each type of fault. For a typical "no fault" Category 5 cable plugged into a 100BASE-T port, 'Good' will be incremented by two after every run of cable diagnostics, once for each cable pair used by a 100BASE-T port. Note that for a 1000BASE-T port, four cable pairs will be tested and so 'Good' will be incremented by four after every successful run.

For a fault condition, an estimated distance to the fault will be calculated and recorded in the system log. For detailed information about which cable pair has been detected to have experienced which type of fault and the corresponding distance to the fault, please refer to the system log file.

Note

The "Runs" parameter cannot be changed while cable diagnostics are running on a port. In order to change the value, stop the diagnostic run on the port, change the "Runs" parameter, and restart diagnostics.

On ports that do not support cable diagnostics, "N/A" will be shown as the cable diagnostics state and any settings made to the "Runs" and "Calibration" fields will be discarded.

2.2.4.3. Calibrating Estimated Distance To Fault

Take the following steps to calibrate the "Calib" parameter (the estimated distance to fault):

1. Pick a particular port for which calibration is needed.
2. Connect an Ethernet cable with a known length (e.g. 50m) to the port.
3. Do not connect the other end of the cable to any link partner.
4. Run cable diagnostics a few times on the port. OPEN fault should be detected.
5. Find the average distance to the OPEN fault recorded in the log and compare it to the known length of the cable. The difference can be used as the calibration value.
6. Enter the calibration value and run cable diagnostics a few more times.
7. The distance to the OPEN fault should now be at a similar distance to the actual cable length.
8. The distance to the fault for the selected port is now calibrated.

2.2.5. Link Detection Options

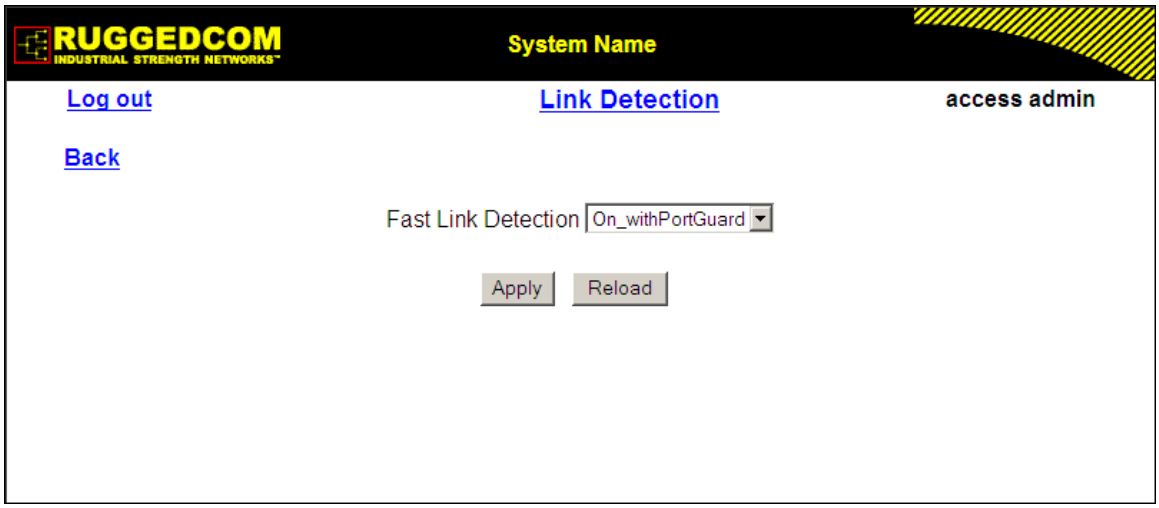

Figure 2.10. Link Detection Form

Fast Link Detection

```
Synopsis: { Off, On, On_withPortGuard }
Default: On_withPortGuard
```

This parameter provides system protection against a faulty end device generating an improper link integrity signal. When a faulty end device or a mismatched fiber port is connected to the unit, a large number of continuous link state changes can be reported in a short period of time. This high rate of link state changes can render the system unresponsive.

Three different settings are available for this parameter:

- **ON_withPortGuard** - This is the recommended setting. With this setting, an extended period (>two minutes) of excessive link state changes reported by a port will prompt the Port Guard feature to disable Fast Link Detection on that port permanently and to raise an alarm. By disabling Fast Link Detection on the port, excessive link state changes can no longer consume a substantial amount of system resources. Note, however, that if Fast Link Detection is disabled, the port will need a longer time to detect a link failure. If the port is part of a spanning tree, this could result in a longer network recovery time, of up to two seconds. Once Port Guard has disabled Fast Link Detection on a particular port, the user can re-enable it by clearing the alarm.

- **ON** - In certain special cases where prolonged, frequent link state change constitutes legitimate link operation, using this setting prevents the system from disabling Fast Link Detection on the port in question. If excessive link state changes persist for more than two minutes on a particular port, an alarm will be generated to warn about the observed bouncing link. If the condition of excessive link state changes is resolved later on, the alarm will be cleared automatically. Since this option does not disable Fast Link Detection, a persistent bouncing link could affect the response time of the system. This setting should be used with caution.
- **OFF** - Turning this parameter OFF will disable Fast Link Detection completely. The switch will need a longer time to detect a link failure. This will result in a longer network recovery time of up to two seconds. This option should only be used if fast link failure detection is not needed.

Note

When Fast Link Detection is enabled, the system prevents link state change processing from consuming all available CPU resources. If Port Guard is not used, however, it is possible for almost all available CPU time to be consumed by frequent link state changes, which could have a negative impact on overall system responsiveness.

2.2.6. Port Status

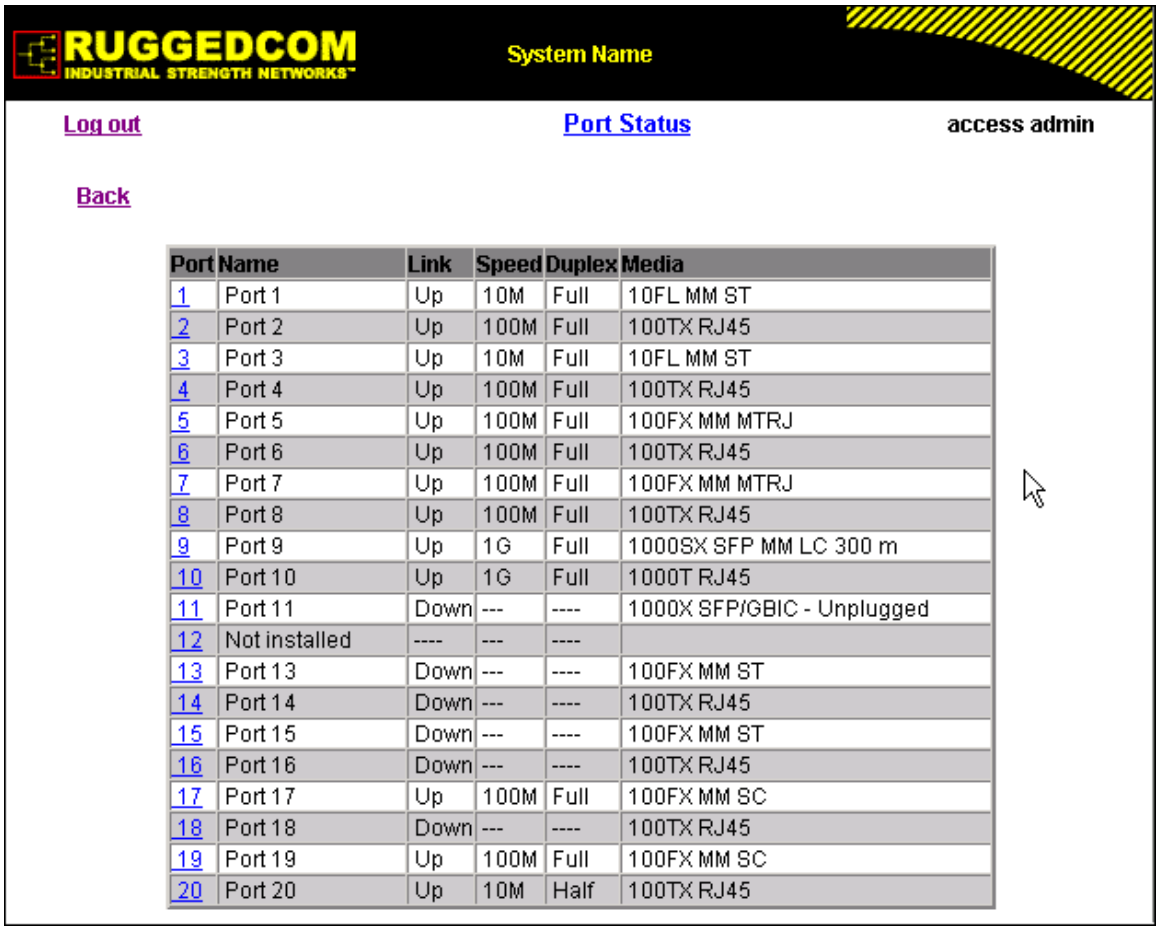

Port	Name	Link	Speed	Duplex	Media
1	Port 1	Up	10M	Full	10FL MM ST
2	Port 2	Up	100M	Full	100TX RJ45
3	Port 3	Up	10M	Full	10FL MM ST
4	Port 4	Up	100M	Full	100TX RJ45
5	Port 5	Up	100M	Full	100FX MM MTRJ
6	Port 6	Up	100M	Full	100TX RJ45
7	Port 7	Up	100M	Full	100FX MM MTRJ
8	Port 8	Up	100M	Full	100TX RJ45
9	Port 9	Up	1G	Full	1000SX SFP MM LC 300 m
10	Port 10	Up	1G	Full	1000T RJ45
11	Port 11	Down	---	----	1000X SFP/GBIC - Unplugged
12	Not installed	----	---	----	
13	Port 13	Down	---	----	100FX MM ST
14	Port 14	Down	---	----	100TX RJ45
15	Port 15	Down	---	----	100FX MM ST
16	Port 16	Down	---	----	100TX RJ45
17	Port 17	Up	100M	Full	100FX MM SC
18	Port 18	Down	---	----	100TX RJ45
19	Port 19	Up	100M	Full	100FX MM SC
20	Port 20	Up	10M	Half	100TX RJ45

Figure 2.11. Port Status Table

Port

Synopsis: 1 to maximum port number

The port for which status is provided.

Name

Synopsis: Any 15 characters

A descriptive name that may be used to identify the device connected to that port.

Link

Synopsis: { ----, ----, Down, Up }

The port's link status.

Speed

Synopsis: { ---, 10, 100, 1000 }

The port's current speed.

Duplex

Synopsis: { ----, Half, Full }

The port's current duplex status.

Media Type

Synopsis: Any 31 characters

Provides user with the description of installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, media description is displayed as per SFF-8472 specification, if transceiver is plugged into the module, e.g., 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST, 1000SX SFP LC S SL M5

2.2.7. Resetting Ports

This command performs a reset of the specified Ethernet ports. This action is useful for forcing re-negotiation of speed and duplex mode or in situations where the link partner has latched into an inappropriate state.

2.3. Troubleshooting

Problem One

One of my links seems to be fine at low traffic levels, but starts to fail as traffic rates increase.

One of my links pings OK but has problems with FTP/SQL/HTTP/…

A possible cause of intermittent operation is that of a ‘duplex mismatch’. If one end of the link is fixed to full-duplex and the peer auto-negotiates, the auto-negotiating end falls back to half-duplex operation. At lower traffic volumes, the link may display few if any errors. As the traffic volume rises, the fixed negotiation side will begin to experience dropped packets while the auto-negotiating side will experience collisions. Ultimately, as traffic loads approach 100%, the link will become entirely unusable.

Note

The ping command with flood options is a useful tool for testing commissioned links. The command "ping 192.168.0.1 500 2" can be used to issue 500 pings each separated by two milliseconds to the next switch. If the link used is of high quality, then no pings should be lost and the average round trip time should be small.

Problem Two

I am trying to use the LFI protection feature but my links won't even come up.

Is it possible that the peer also has LFI enabled? If both sides of the link have LFI enabled, then both sides will withhold link signal generation from each other.

3. Ethernet Statistics

ROS® Ethernet Statistics provide you with the following abilities:

- Viewing basic Ethernet statistics.
- Viewing and clearing detailed Ethernet statistics.
- Configuring RMON History control.
- Viewing collected RMON History samples.
- Configuring RMON Alarms.
- Configuring RMON Events.
- Viewing collected RMON Event logs.

The Ethernet Statistics menu is accessible from the main menu.

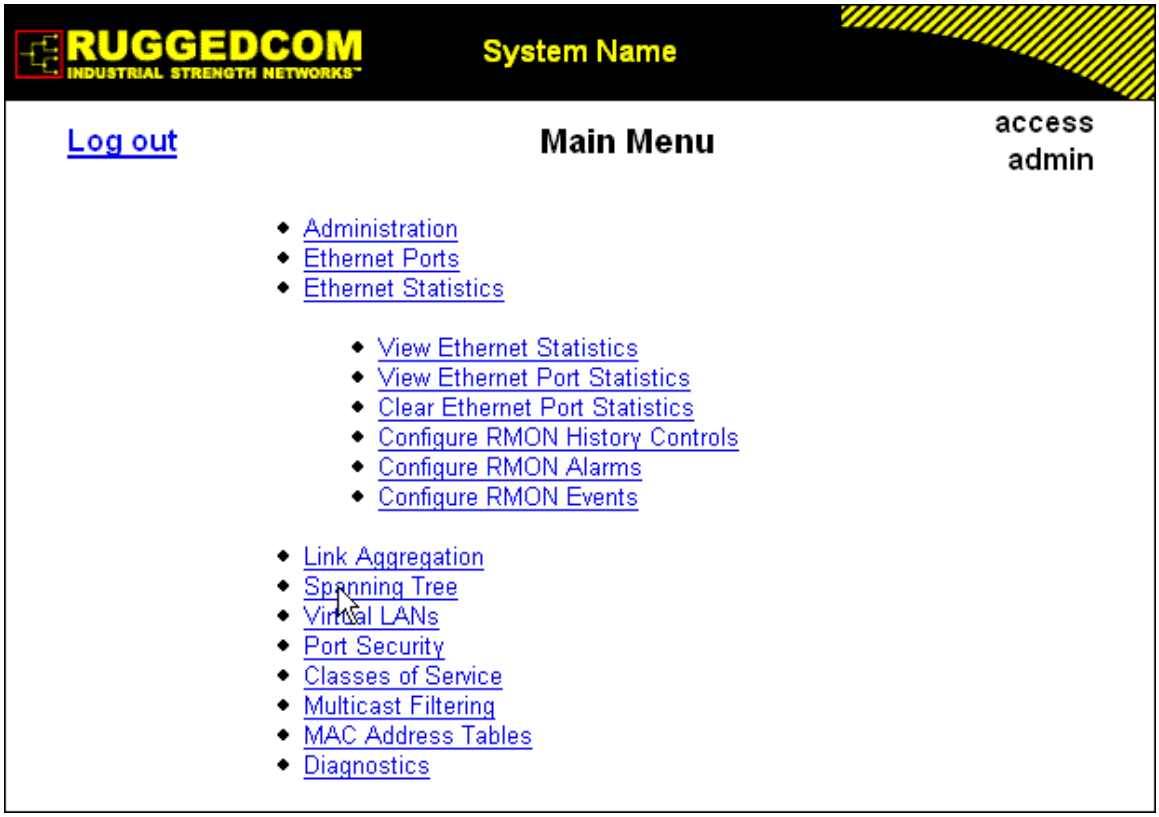

Figure 3.1. Ethernet Port Statistics Menu

3.1. Viewing Ethernet Statistics

This table provides basic Ethernet statistics information which is reset periodically, every few seconds. This traffic view is useful when the origin and destination of a traffic flow need to be determined.

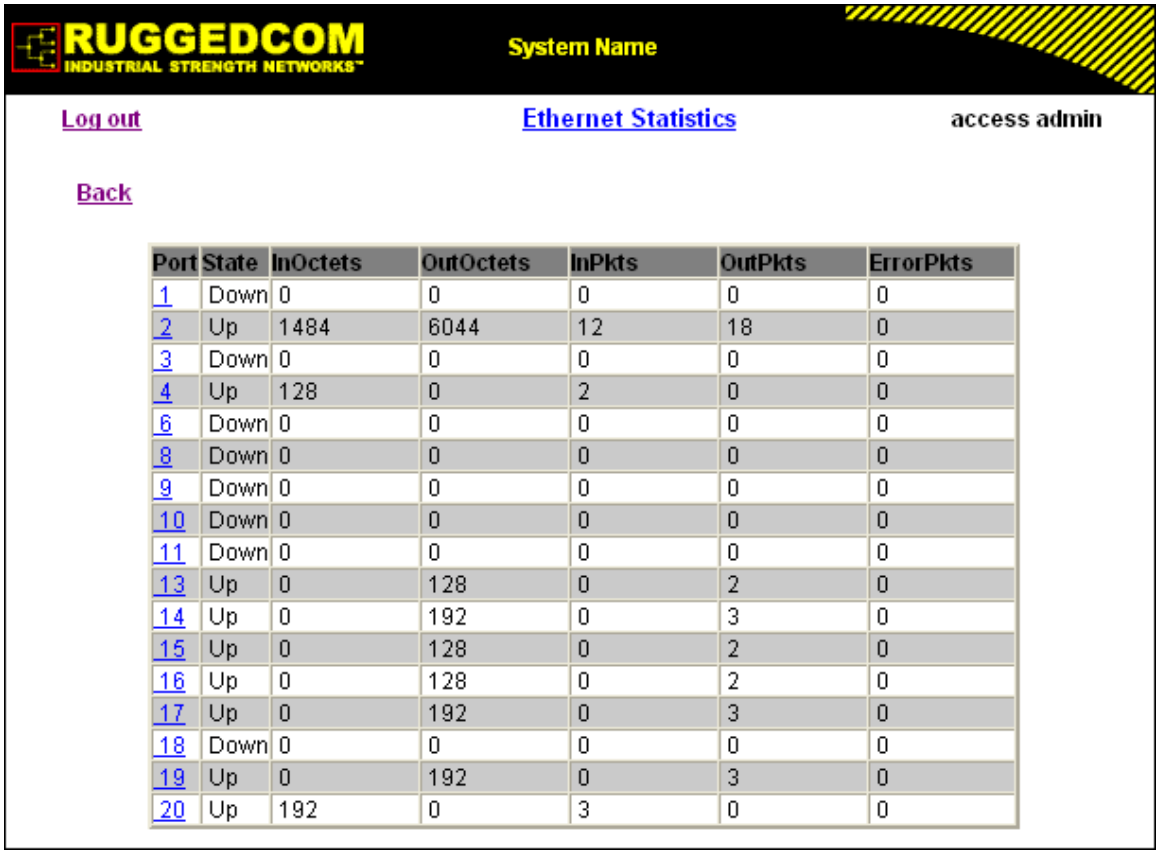

Port	State	InOctets	OutOctets	InPkts	OutPkts	ErrorPkts
1	Down	0	0	0	0	0
2	Up	1484	6044	12	18	0
3	Down	0	0	0	0	0
4	Up	128	0	2	0	0
6	Down	0	0	0	0	0
8	Down	0	0	0	0	0
9	Down	0	0	0	0	0
10	Down	0	0	0	0	0
11	Down	0	0	0	0	0
13	Up	0	128	0	2	0
14	Up	0	192	0	3	0
15	Up	0	128	0	2	0
16	Up	0	128	0	2	0
17	Up	0	192	0	3	0
18	Down	0	0	0	0	0
19	Up	0	192	0	3	0
20	Up	192	0	3	0	0

Figure 3.2. Ethernet Statistics Table

Figure 2: Ethernet Statistics Table

Port

```
Synopsis: 1 to maximum port number
```

The port number as seen on the front plate silkscreen of the switch.

State

```
Synopsis: { ----, Down, Up }
```

The port link status.

InOctets

```
Synopsis: 0 to 4294967295
```

The number of octets in received good packets (Unicast+Multicast+Broadcast) and dropped packets.

OutOctets

```
Synopsis: 0 to 4294967295
```

The number of octets in transmitted good packets.

InPkts

```
Synopsis: 0 to 4294967295
```

The number of received good packets (Unicast+Multicast+Broadcast) and dropped packets.

OutPkts

```
Synopsis: 0 to 4294967295
```

The number of transmitted good packets.

ErrorPkts

`Synopsis: 0 to 4294967295`

The number of any type of erroneous packet.

3.2. Viewing Ethernet Port Statistics

Ethernet port statistics provide a detailed view of the traffic. This is useful when the exact source of error or traffic mix needs to be determined.

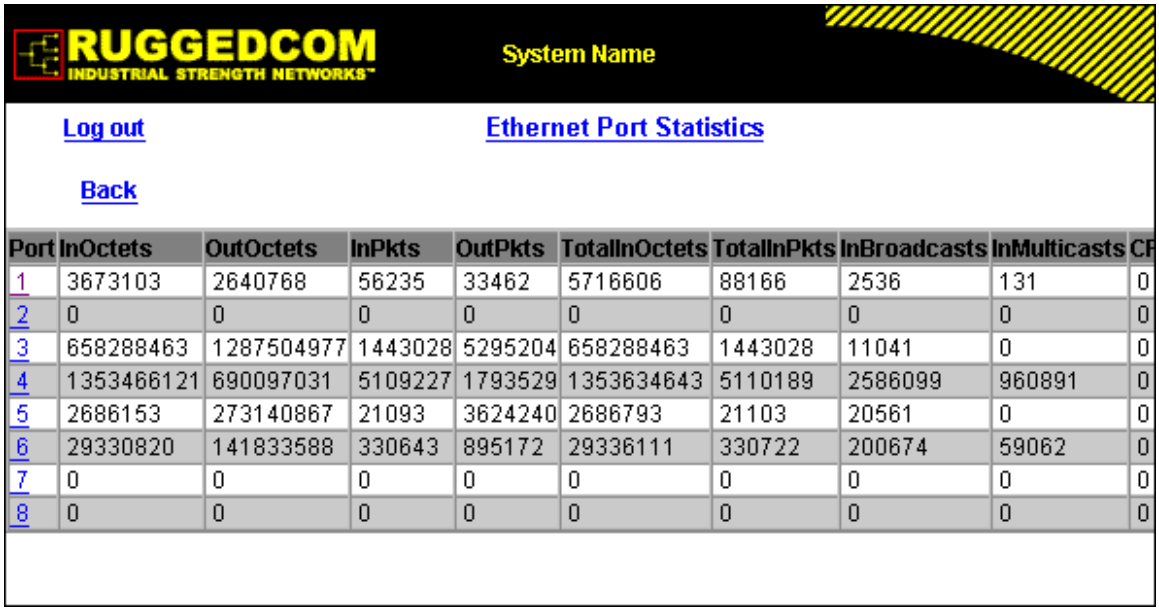

Port	InOctets	OutOctets	InPkts	OutPkts	TotalInOctets	TotalInPkts	InBroadcasts	InMulticasts	CF
1	3673103	2640768	56235	33462	5716606	88166	2536	131	0
2	0	0	0	0	0	0	0	0	0
3	658288463	1287504977	1443028	5295204	658288463	1443028	11041	0	0
4	1353466121	690097031	5109227	1793529	1353634643	5110189	2586099	960891	0
5	2686153	273140867	21093	3624240	2686793	21103	20561	0	0
6	29330820	141833588	330643	895172	29336111	330722	200674	59062	0
7	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0

Figure 3.3. Ethernet Port Statistics Table

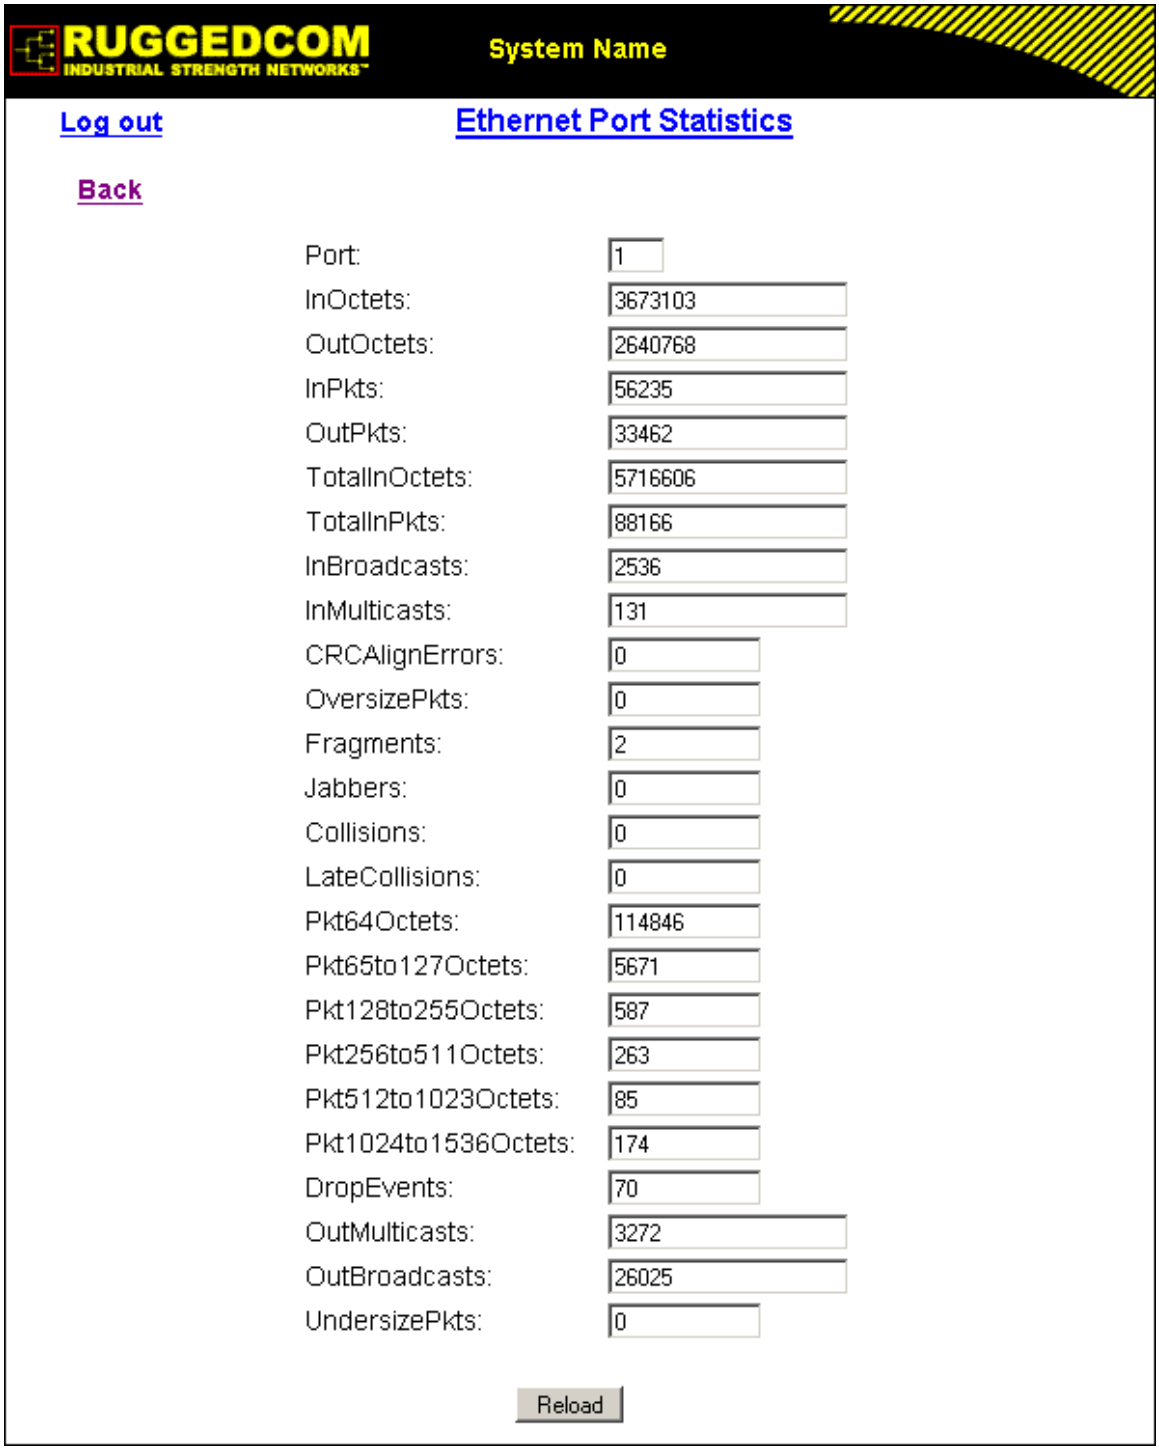

Figure 3.4. Ethernet Port Statistics Form

Port

```
Synopsis: 1 to maximum port number
```

The port number as seen on the front plate silkscreen of the switch.

InOctets

```
Synopsis: 0 to 18446744073709551615
```

The number of octets in both received packets (Unicast+Multicast+Broadcast) and dropped packets.

OutOctets

`Synopsis: 0 to 18446744073709551615`

The number of octets in transmitted packets.

InPkts

`Synopsis: 0 to 18446744073709551615`

The number of received good packets (Unicast+Multicast+Broadcast) and dropped packets.

OutPkts

`Synopsis: 0 to 18446744073709551615`

The number of transmitted good packets.

TotalInOctets

`Synopsis: 0 to 18446744073709551615`

The total number of octets of all received packets. This includes data octets of rejected and local packets which are not forwarded to the switching core for transmission. It should reflect all the data octets received on the line.

TotalInPkts

`Synopsis: 0 to 18446744073709551615`

The number of received packets. This includes rejected, dropped local, and packets which are not forwarded to the switching core for transmission. It should reflect all packets received on the line.

InBroadcasts

`Synopsis: 0 to 18446744073709551615`

The number of Broadcast packets received.

InMulticasts

`Synopsis: 0 to 18446744073709551615`

The number of Multicast packets received.

CRCAlignErrors

`Synopsis: 0 to 4294967295`

The number of packets received which meet all the following conditions:

1. Packet data length is between 64 and 1536 octets inclusive.
2. Packet has invalid CRC.
3. Collision Event has not been detected.
4. Late Collision Event has not been detected.

OversizePkts

`Synopsis: 0 to 4294967295`

The number of packets received with data length greater than 1536 octets and valid CRC.

Fragments

`Synopsis: 0 to 4294967295`

The number of packets received which meet all the following conditions:

1. Packet data length is less than 64 octets.
2. Collision Event has not been detected.

3. Late Collision Event has not been detected.
4. Packet has invalid CRC.

Jabbers

`Synopsis: 0 to 4294967295`

The number of packets which meet all the following conditions:

1. Packet data length is greater that 1536 octets.
2. Packet has invalid CRC.

Collisions

`Synopsis: 0 to 4294967295`

The number of received packets for which Collision Event has been detected.

LateCollisions

`Synopsis: 0 to 4294967295`

The number of received packets for which Late Collision Event has been detected.

Pkt64Octets

`Synopsis: 0 to 4294967295`

The number of received and transmitted packets with size of 64 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets.

Pkt65to127Octets

`Synopsis: 0 to 4294967295`

The number of received and transmitted packets with a size of 65 to 127 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets.

Pkt128to255Octets

`Synopsis: 0 to 4294967295`

The number of received and transmitted packets with a size of 128 to 257 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets.

Pkt256to511Octets

`Synopsis: 0 to 4294967295`

The number of received and transmitted packets with a size of 256 to 511 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets.

Pkt512to1023Octets

`Synopsis: 0 to 4294967295`

The number of received and transmitted packets with a size of 512 to 1023 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets.

Pkt1024to1536Octets

`Synopsis: 0 to 4294967295`

The number of received and transmitted packets with a size of 1024 to 1536 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets.

DropEvents

`Synopsis: 0 to 4294967295`

The number of received packets that are dropped due to lack of receive buffers.

OutMulticasts

`Synopsis: 0 to 18446744073709551615`

The number of transmitted multicast packets. This does not include broadcast packets.

OutBroadcasts

`Synopsis: 0 to 18446744073709551615`

The number of transmitted broadcast packets.

UndersizePkts

`Synopsis: 0 to 18446744073709551615`

The number of received packets which meet all the following conditions:

1. Packet data length is less than 64 octets.
2. Collision Event has not been detected.
3. Late Collision Event has not been detected.
4. Packet has valid CRC.

OutUcastPkts

`Synopsis: 0 to 18446744073709551615`

The number of transmitted unicast packets.

3.3. Clearing Ethernet Port Statistics

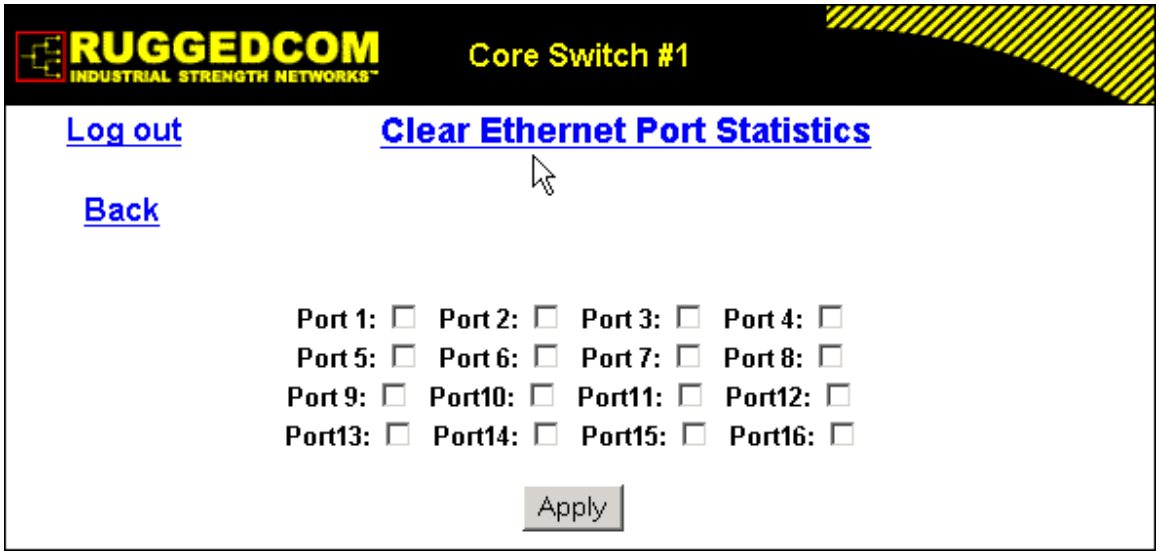

Figure 3.5. Clear Ethernet Port Statistics Form

This command clears Ethernet ports statistics for one or more Ethernet ports. Ports are chosen by checking the corresponding boxes.

3.4. Remote Monitoring (RMON)

The RuggedSwitch® Remote Monitor (RMON) package provides the following capabilities:

- The ability to collect and view historical statistics in order to review performance and operation of Ethernet ports.
- The ability to record a log entry and/or generate an SNMP trap when the rate of occurrence of a specified event is exceeded.

3.4.1. RMON History Controls

The RMON History Controls table programs the switch to take samples of the RMON-MIB history statistics of an Ethernet port at regular intervals.

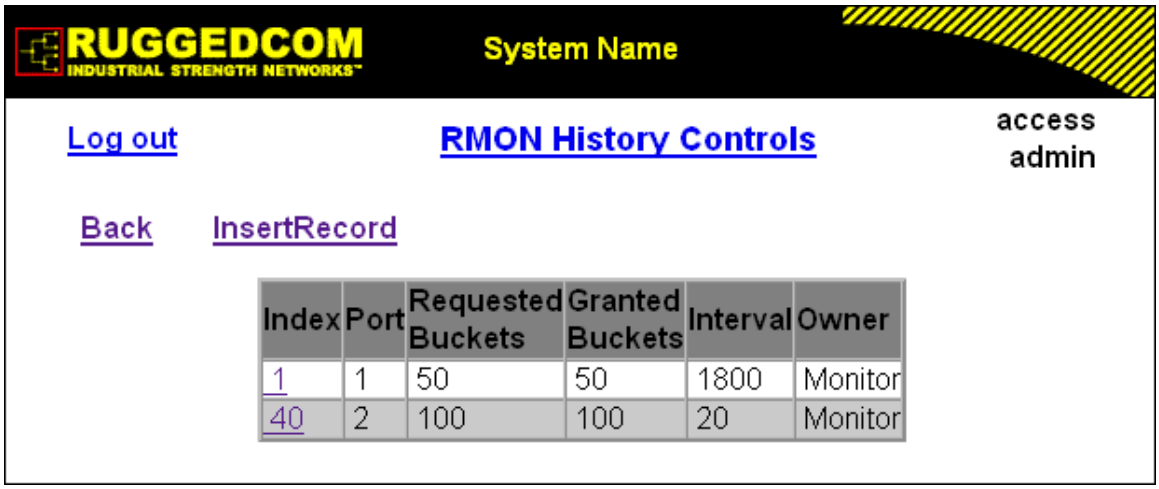

Figure 3.6. RMON History Controls Table

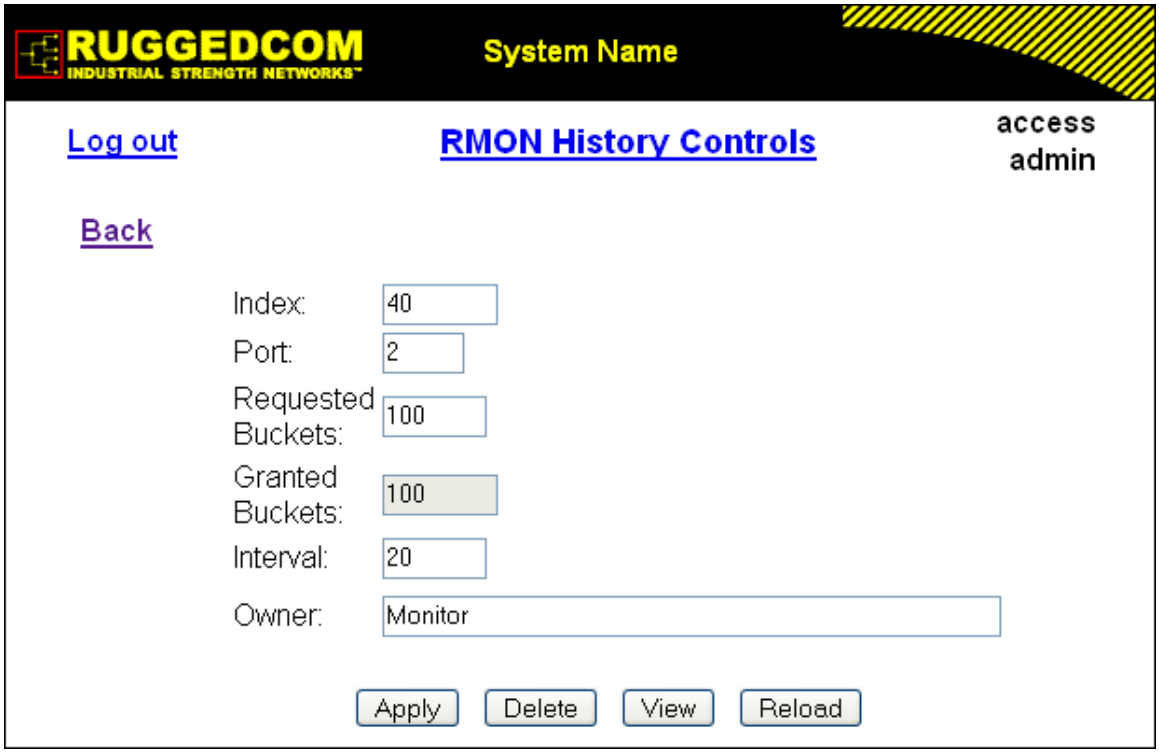

Figure 3.7. RMON History Controls Form

Index

```
Synopsis: 1 to 65535
Default: 1
```

The index of this RMON History Control record.

Port

```
Synopsis: 1 to maximum port number
Default: 1
```

The port number as seen on the front plate silkscreen of the switch.

Requested Buckets

```
Synopsis: 1 to 4000
Default: 50
```

The maximum number of buckets requested for this RMON collection history group of statistics. The range is 1 to 4000. The default is 50.

Granted Buckets

```
Synopsis: 0 to 65535
```

The number of buckets granted for this RMON collection history. This field is not editable.

Interval

```
Synopsis: 1 to 3600
Default: 1800
```

The number of seconds in over which the data is sampled for each bucket. The range is 1 to 3600. The default is 1800.

Owner

```
Synopsis: Any 127 characters
Default: Monitor
```

The owner of this record. It is suggested to start this string with the word 'monitor'.

3.4.2. RMON History Samples

History samples for a particular record in the RMON History Control Table are displayed by selecting a particular record and view option. The index of the record will be included in the resulting menu title of the sample screen.

The table will present a series of samples. The sample number starts with one and increases by one with each new log entry. The oldest samples are deleted in favor of new samples when the allotted buckets are used.

The StartTime field provides the system time when the measurement interval started. The remaining fields provide the counts for each statistic as measured in the sample period.

Statistics collection begins whenever the History Control record is created and when the switch is initialized. As new samples are added, the window is automatically updated.

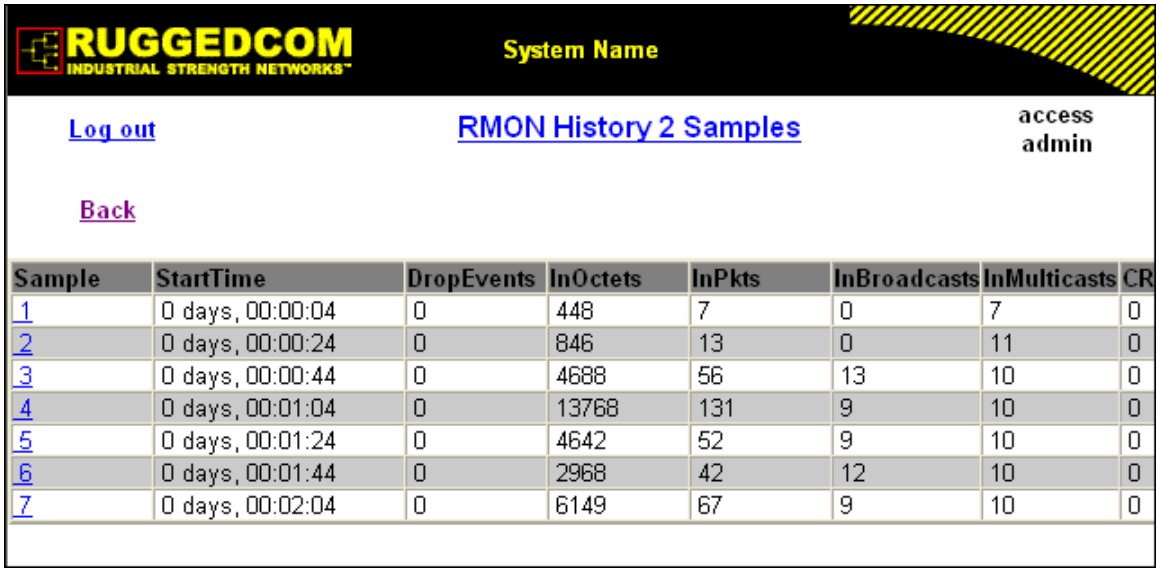

Sample	StartTime	DropEvents	InOctets	InPkts	InBroadcasts	InMulticasts	CR
1	0 days, 00:00:04	0	448	7	0	7	0
2	0 days, 00:00:24	0	846	13	0	11	0
3	0 days, 00:00:44	0	4688	56	13	10	0
4	0 days, 00:01:04	0	13768	131	9	10	0
5	0 days, 00:01:24	0	4642	52	9	10	0
6	0 days, 00:01:44	0	2968	42	12	10	0
7	0 days, 00:02:04	0	6149	67	9	10	0

Figure 3.8. RMON History Samples Table

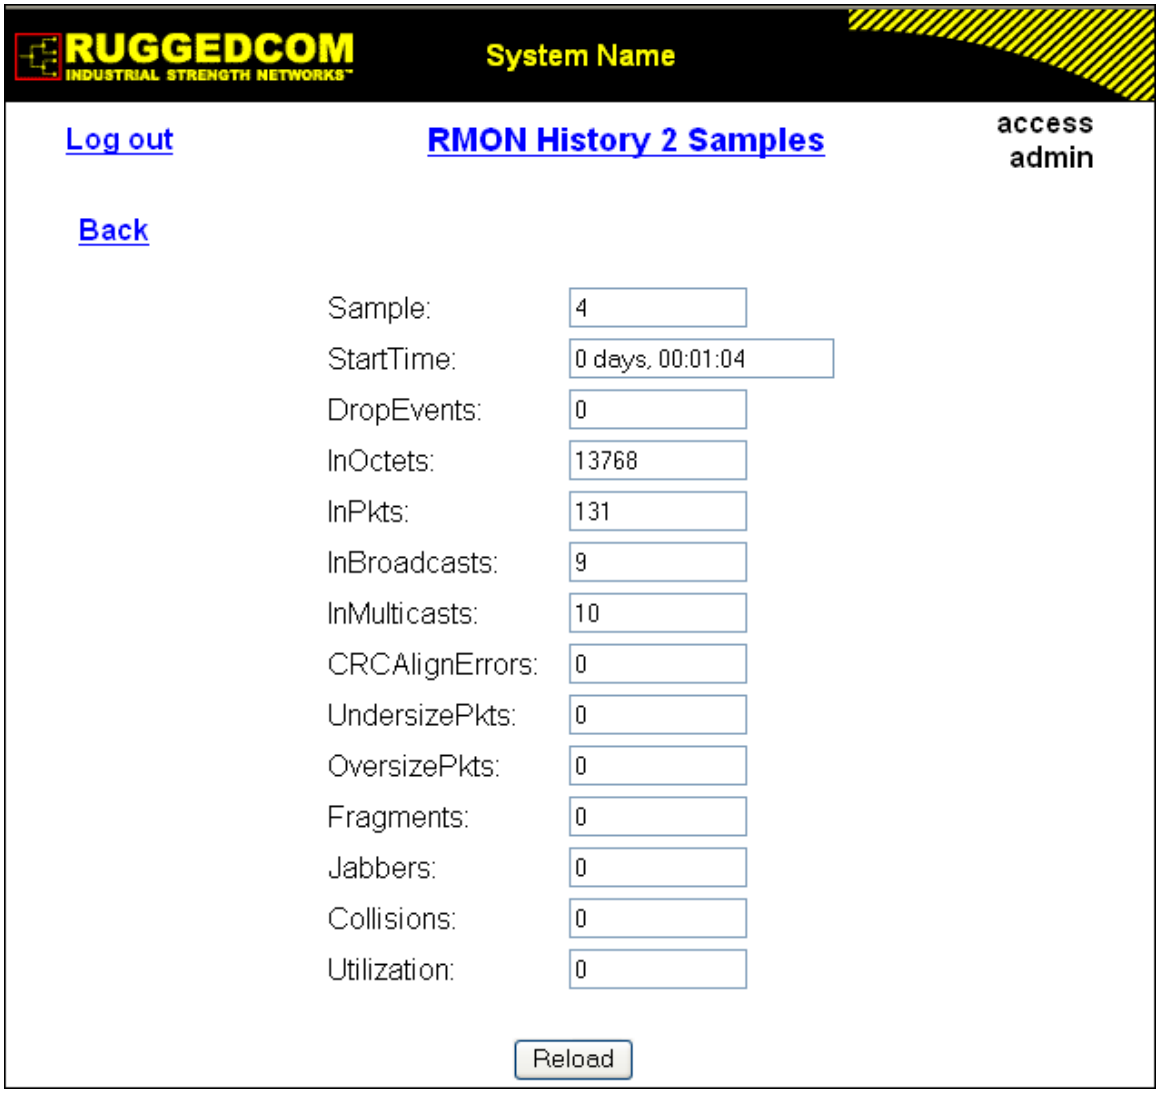

Figure 3.9. RMON History Samples Form

Sample

`Synopsis: 0 to 4294967295`

The sample number taken for this history record.

StartTime

`Synopsis: DDDD days, HH:MM:SS`

The system elapsed time when started interval over which this sample was measured

DropEvents

`Synopsis: 0 to 4294967295`

The number of received packets that are dropped due to lack of receive buffers.

InOctets

`Synopsis: 0 to 4294967295`

The number of octets in good packets (Unicast+Multicast+Broadcast) and dropped packets received.

InPkts

`Synopsis: 0 to 4294967295`

The number of good packets (Unicast+Multicast+Broadcast) and dropped packets received.

InBroadcasts

`Synopsis: 0 to 4294967295`

The number of broadcast packets received.

InMulticasts

`Synopsis: 0 to 4294967295`

The number of multicast packets received.

CRCAlignErrors

`Synopsis: 0 to 4294967295`

The number of packets received that meet all the following conditions:

1. Packet data length is between 64 and 1536 octets inclusive.
2. Packet has invalid CRC.
3. Collision Event has not been detected.
4. Late Collision Event has not been detected.

UndersizePkts

`Synopsis: 0 to 4294967295`

The number of received packets that meet all the following conditions:

1. Packet data length is less than 64 octets.
2. Collision Event has not been detected.
3. Late Collision Event has not been detected.
4. Packet has valid CRC.

OversizePkts

`Synopsis: 0 to 4294967295`

The number of packets received with data length greater than 1536 octets and valid CRC.

Fragments

`Synopsis: 0 to 4294967295`

The number of packets received that meet all the following conditions:

1. Packet data length is less than 64 octets.
2. Collision Event has not been detected.
3. Late Collision Event has not been detected.
4. Packet has invalid CRC.

Jabbers

`Synopsis: 0 to 4294967295`

The number of packets that meet all the following conditions:

1. Packet data length is greater that 1536 octets.
2. Packet has invalid CRC.

Collisions

`Synopsis: 0 to 4294967295`

The number of received packets for which Collision Event has been detected.

Utilization

```
Synopsis: 0 to 100
```

The best estimate of the mean physical layer network utilization on this interface during this sampling interval (in percent).

3.4.3. RMON Alarms

The RMON Alarm table configures the switch to examine the state of a specific statistical variable.

The record of this table contains an upper and a lower threshold for legal values of the statistic in a given interval. This provides the ability to detect events occurring more quickly than a specified maximum rate or less quickly than a specified minimum rate.

When a statistic value's rate of change exceeds its limits, an internal alarm of INFO level is always generated. Internal alarms can be viewed using the Diagnostics menu, View Alarms command.

Additionally, a statistic threshold crossing can result in further activity. The RMON Alarm record can be configured to point to a particular RMON Event Record, which can generate an SNMP trap, an entry in the switch's event log or both. The RMON Event Record can "steer" alarms towards different users defined in SNMP Users table.

The alarm record can point to a different event record for each of the thresholds, so combinations such as "trap on rising threshold" or "trap on rising threshold, log and trap on falling threshold" are possible.

Each RMON alarm may be configured such that its first instance occurs only for rising, falling, or all threshold excessions.

The ability to configure upper and lower thresholds on the value of a measured statistic provides for the ability to add hysteresis to the alarm generation process.

If the value of the measured statistic over time is compared to a single threshold, alarms will be generated each time the statistic crosses the threshold. If the statistic's value fluctuates around the threshold, an alarm can be generated every measurement period. Programming different upper and lower thresholds eliminates spurious alarms. The statistic value must "travel" between the thresholds before alarms can be generated.

The following figure illustrates the very different patterns of alarm generation resulting from a statistic sample and the same sample with hysteresis applied.

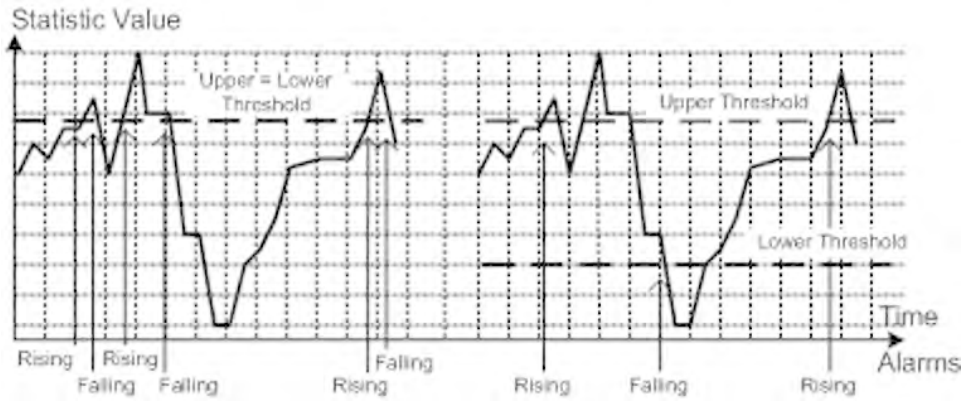

Figure 3.10. The Alarm Process

There are two methods to evaluate a statistic in order to determine when to generate an event; these are the delta and absolute methods.

For most statistics, such as line errors, it is appropriate to alarm when a rate is exceeded. The alarm record defaults to the “delta” measurement method, which examines changes in a statistic at the end of each measurement period.

It may be desirable to alarm when the total, or absolute, number of events crosses a threshold. In this case, set the measurement period type to “absolute”.

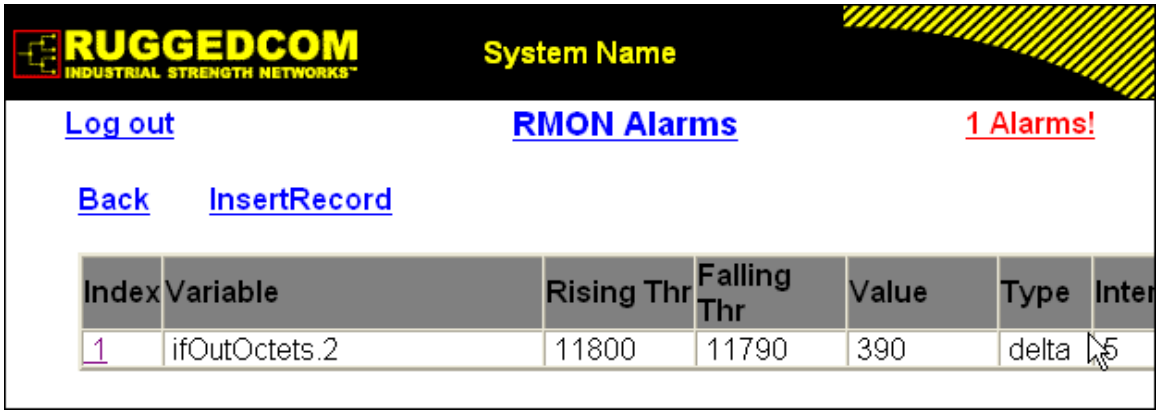

Figure 3.11. RMON Alarms Table

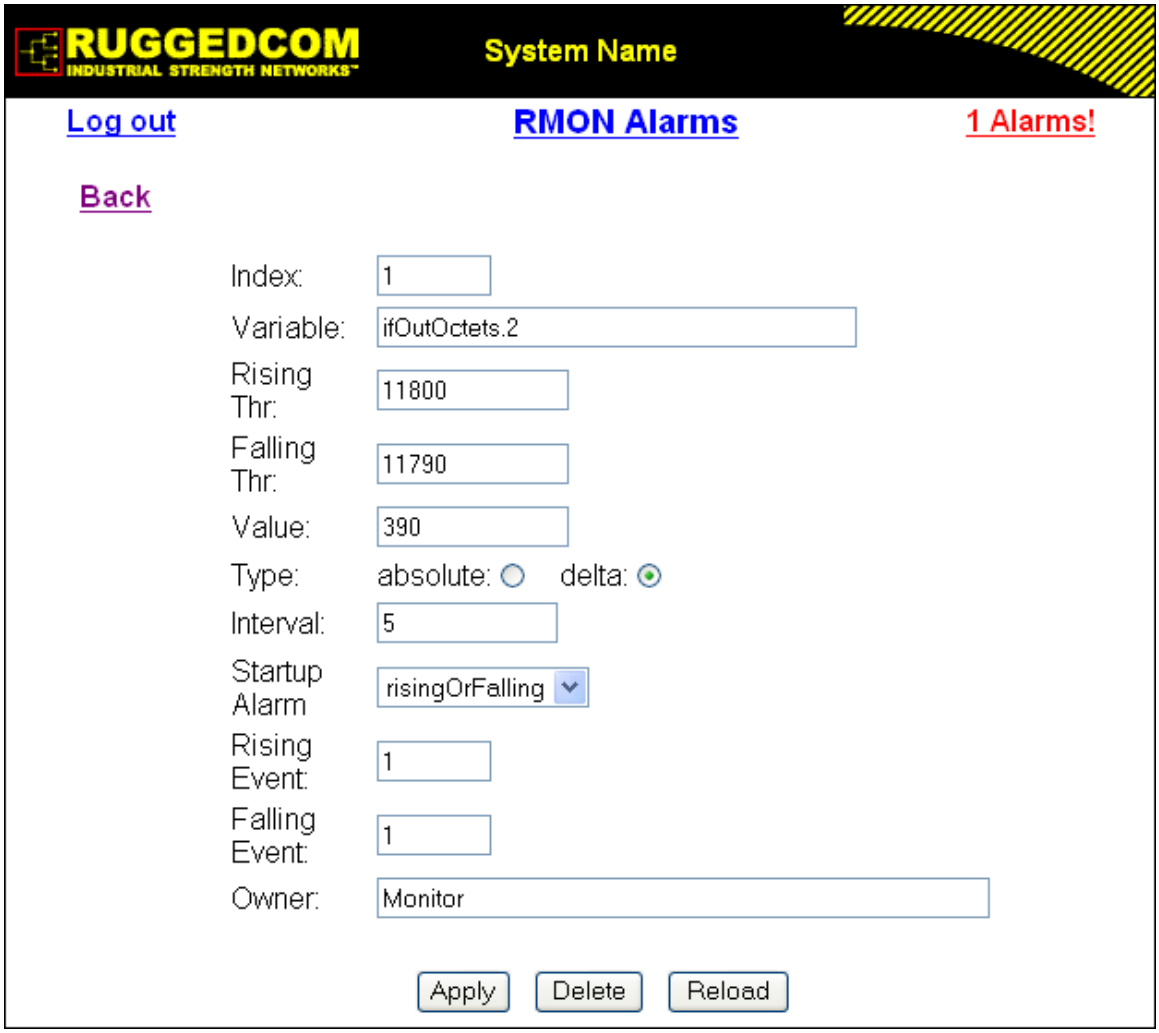

Figure 3.12. RMON Alarms Form

Index

```
Synopsis: 1 to 65535
Default: 2
```

The index of this RMON Alarm record.

Variable

```
Synopsis: SNMP Object Identifier - up to 39 characters
Default: ifOutOctets.2
```

The SNMP object identifier (OID) of the particular variable to be sampled. Only variables that resolve to an ASN.1 primitive type INTEGER (INTEGER, Integer32,Counter32, Counter64, Gauge, or TimeTicks) may be sampled. A list of objects can be printed using shell command 'rmon'. The OID format: objectName.index1.index2… where index format depends on index object type.

Rising Thr

```
Synopsis: 0 to 2147483647
Default: 11800
```

A threshold for the sampled variable. When the current sampled variable value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event will be generated. A single event will also be generated if the first sample created after this record is greater than or equal to this threshold and the associated startup alarm is equal to 'rising'. After a rising alarm is generated, another such event will not be generated until the sampled value falls below this threshold and reaches the value of FallingThreshold.

Falling Thr

```
Synopsis: 0 to 2147483647
Default: 11790
```

A threshold for the sampled variable. When the current sampled variable value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event will be generated. A single event will also be generated if the first sample created after this record is less than or equal to this threshold and the associated startup alarm is equal to 'falling'. After a falling alarm is generated, another such event will not be generated until the sampled value rises above this threshold and reaches the value of RisingThreshold.

Value

```
Synopsis: 0 to 2147483647
```

The value of a monitored object during the last sampling period. The presentation of the value depends on the sample type ('absolute' or 'delta').

Type

```
Synopsis: { absolute, delta }
Default: delta
```

The method of sampling the selected variable and calculating the value to be compared against the thresholds. The value of the sample type can be 'absolute' or 'delta'.

Interval

```
Synopsis: 0 to 2147483647
Default: 5
```

The number of seconds during which the data is sampled and compared with the rising and falling thresholds.

Startup Alarm

```
Synopsis: { rising, falling, risingOrFalling }
Default: risingOrFalling
```

The alarm that may be sent when this record is first created if the condition for raising an alarm is met. The value of a startup alarm can be 'rising', 'falling' or 'risingOrFalling'

Rising Event

```
Synopsis: 0 to 65535
Default: 1
```

The index of the event that is used when a falling threshold is crossed. If there is no corresponding entry in the Event Table, then no association exists. In particular, if this value is zero, no associated event will be generated.

Falling Event

```
Synopsis: 0 to 65535
```

```
Default: 1
```

The index of the event that is used when a rising threshold is crossed. If there is no corresponding entry in the Event Table, then no association exists. In particular, if this value is zero, no associated event will be generated.

Owner

```
Synopsis: Any 127 characters
Default: Monitor
```

The owner of this record. It is suggested to start this string with the word 'monitor'.

3.5. RMON Events

The RMON Events Table stores profiles of behavior used in event logging. These profiles are used by RMON Alarm records to send traps and to log events.

Each record may specify that an alarms log entry be created on its behalf whenever the event occurs. Each entry may also specify that a notification should occur by way of SNMP trap messages. In this case, the user for the trap message is given as parameter “Community”.

Two traps are defined: risingAlarm and fallingAlarm.

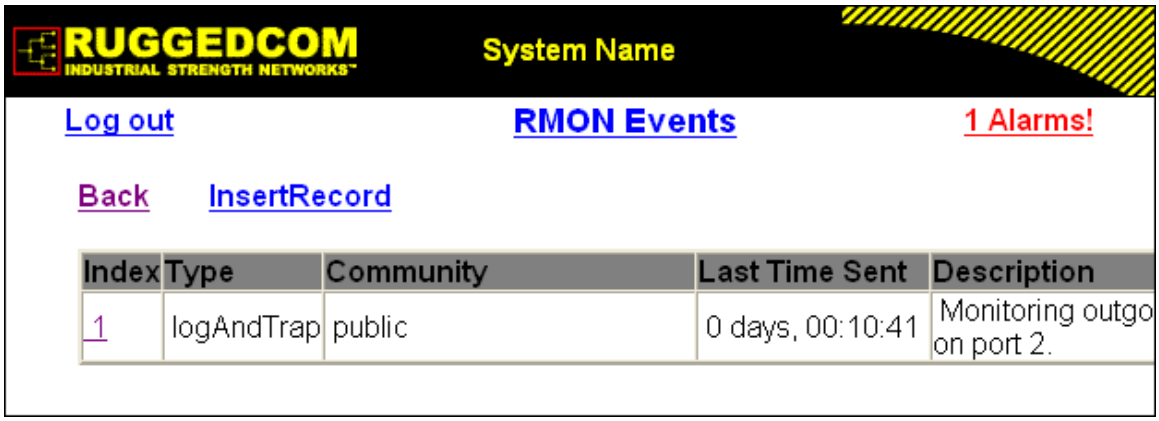

Figure 3.13. RMON Events Table

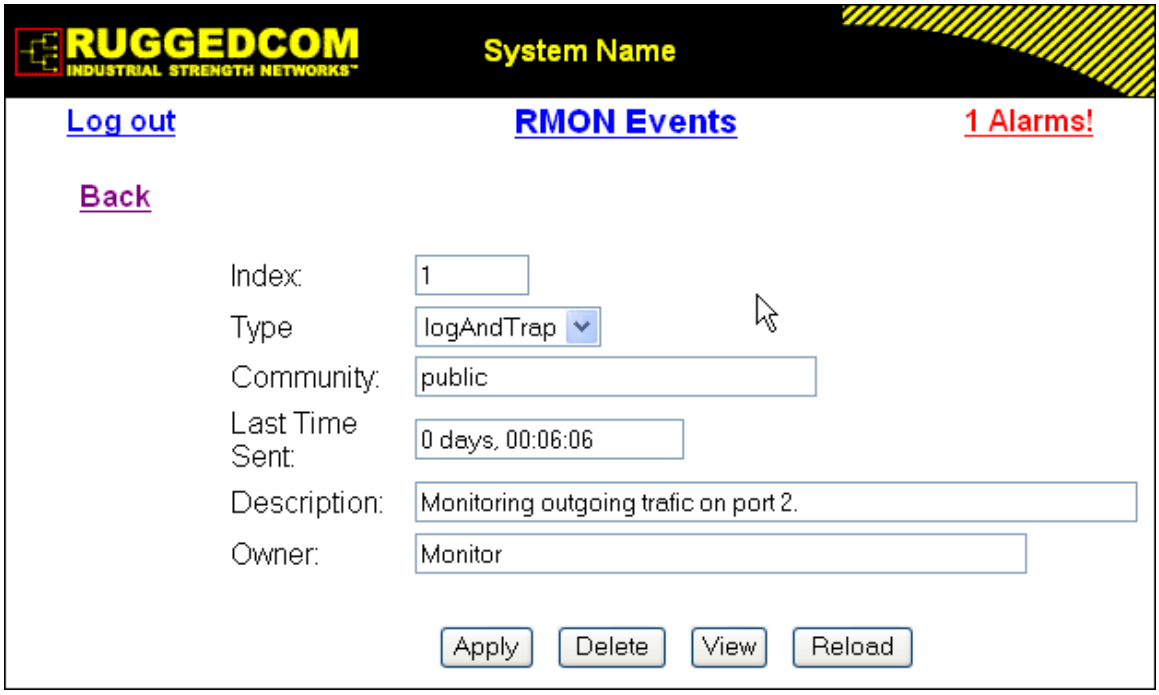

Figure 3.14. RMON Events Form

Index

```
Synopsis: 1 to 65535
Default: 2
```

The index of this RMON Event record.

Type

```
Synopsis: { none, log, snmpTrap, logAndTrap }
Default: logAndTrap
```

The type of notification that the probe will make about this event. In the case of 'log', an entry is made in the RMON Log table for each event. In the case of snmp_trap, an SNMP trap is sent to one or more management stations.

Community

```
Synopsis: Any 31 characters
Default: public
```

If the SNMP trap is to be sent, it will be sent to the SNMP community specified by this string.

Last Time Sent

```
Synopsis: DDDD days, HH:MM:SS
```

The time from last reboot at the time this event entry last generated an event. If this entry has not generated any events, this value will be 0.

Description

```
Synopsis: Any 127 characters
Default: Monitoring outgoing traffic on port 2.
```

A comment describing this event.

Owner

```
Synopsis: Any 127 characters
```

`Default: Monitor`

The owner of this event record. It is suggested to start this string with the word 'monitor'.

3.6. RMON Event Log

Event logs for a particular record in the RMON Events Table can be viewed by selecting a particular record and view option. The index of the record will be included in the resulting menu title of the log table.

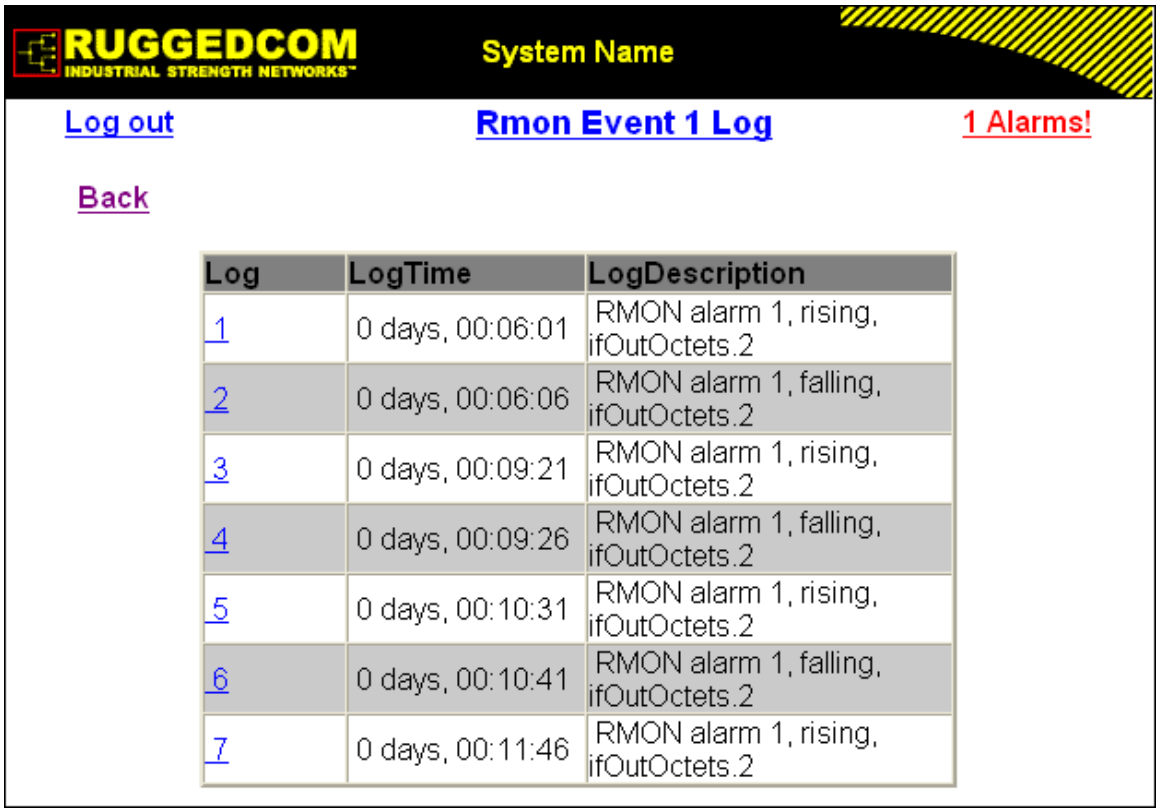

Log	LogTime	LogDescription
1	0 days, 00:06:01	RMON alarm 1, rising, ifOutOctets.2
2	0 days, 00:06:06	RMON alarm 1, falling, ifOutOctets.2
3	0 days, 00:09:21	RMON alarm 1, rising, ifOutOctets.2
4	0 days, 00:09:26	RMON alarm 1, falling, ifOutOctets.2
5	0 days, 00:10:31	RMON alarm 1, rising, ifOutOctets.2
6	0 days, 00:10:41	RMON alarm 1, falling, ifOutOctets.2
7	0 days, 00:11:46	RMON alarm 1, rising, ifOutOctets.2

Figure 3.15. RMON Event Log Table

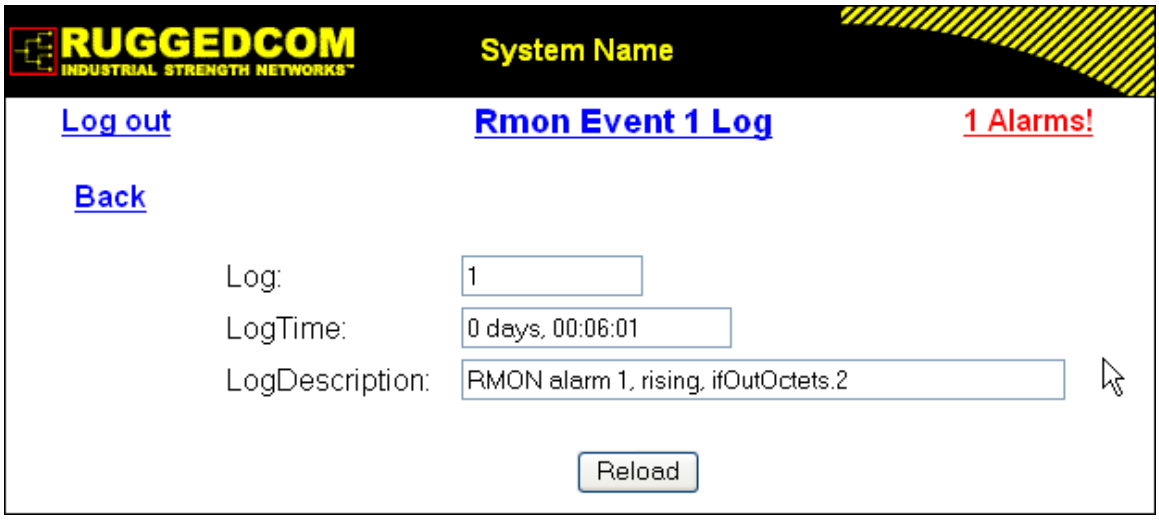

Figure 3.16. RMON Event Log Form

Log

Synopsis: 0 to 4294967295

The index (log) taken for this log record.

LogTime

Synopsis: DDDD days, HH:MM:SS

The system elapsed time when this log was created.

LogDescription

Synopsis: Any 49 characters

The description of the event that activated this log entry.

4. Link Aggregation

Link Aggregation is also known as port trunking or port bundling.

ROS® provides you with the following Link Aggregation features:

- Support for up to 15 port trunks.

Note

The actual maximum number of port trunks depends on the number of ports in the switch (at least two ports are required to compose a port trunk)

- Up to 8 ports can be aggregated in one port trunk.
- Highly randomized load balancing between the aggregated links based on both source and destination MAC addresses of the forwarded frames.

4.1. Link Aggregation Operation

Link Aggregation provides you with the ability to aggregate several Ethernet ports into one logical link (port trunk) with higher bandwidth. Link Aggregation can be used for two purposes:

- To obtain increased, linearly incremental, link bandwidth.
- To improve network reliability by creating link redundancy. If one of the aggregated links fails, the switch will balance the traffic between the remaining links.

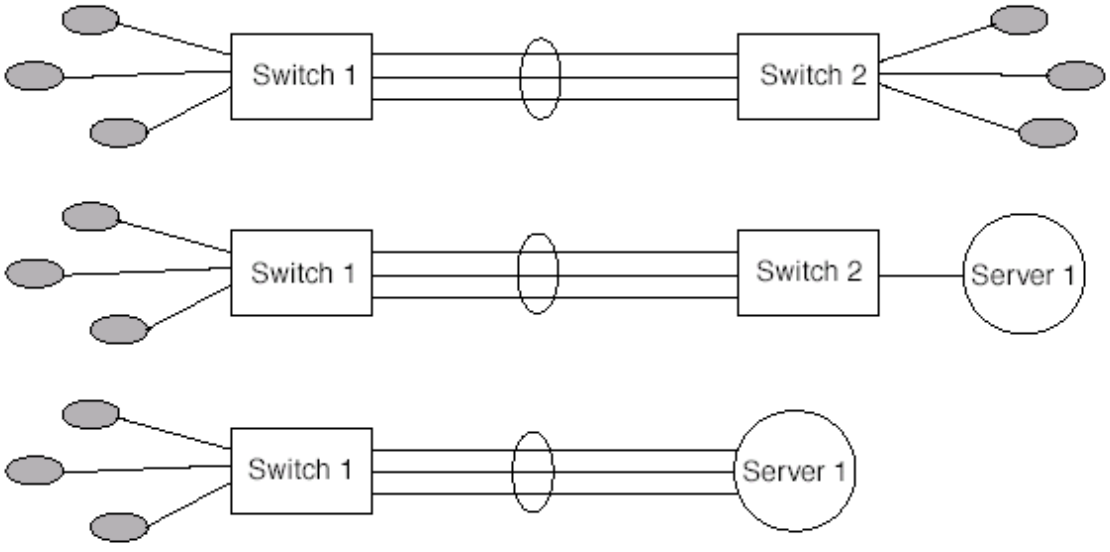

Figure 4.1. Link Aggregation Examples

4.1.1. Link Aggregation Rules

- Any port can belong to only one port trunk at a time.

- The aggregated port with the lowest port number is called the Port Trunk Primary Port. Other ports in the trunk are called Secondary Ports.
- Layer 2 features (e.g. STP, VLAN, CoS, Multicast Filtering) treat a port trunk as a single link.
  - If STP puts an aggregated port in blocking/forwarding, it does it for the whole port trunk
  - If one of the aggregated ports joins/leaves a multicast group (e.g. via IGMP or GMRP), all other ports in the trunk will join/leave too.
  - Any port configuration parameter (e.g. VLAN, CoS) change will be automatically applied to all ports in the trunk.
  - Configuration/status parameters of the secondary ports will not be shown and their port numbers will be simply listed next to the primary port number in the appropriate configuration/status UI sessions. For example:

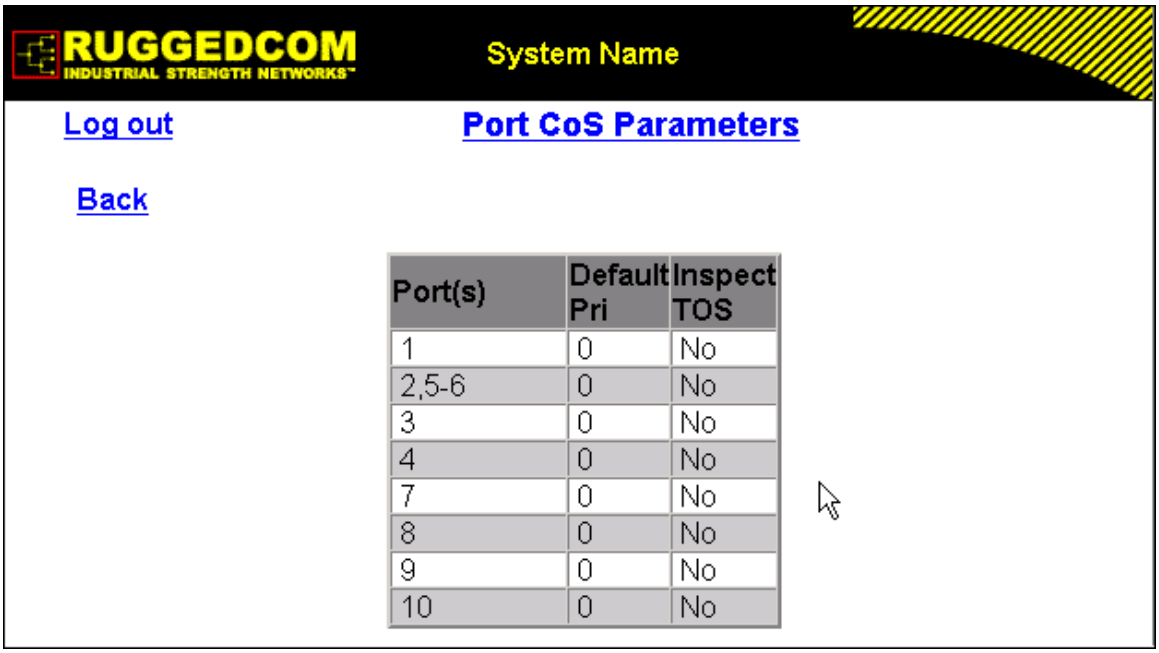

Figure 4.2. Displaying Port Trunk Secondary Ports in Layer 2 Feature Configuration

  - When a secondary port is added to a port trunk, it inherits all the configuration settings of the primary port. When this secondary port is removed from the port trunk, the settings it had previous to the aggregation are restored.
- Physical layer features (e.g. physical link configuration, link status, rate limiting, Ethernet statistics) will still treat each aggregated port separately.
  - Physical configuration/status parameters will NOT be automatically applied to other ports in the trunk and will be displayed for each port as usual.
  - Make sure that only ports with the same speed and duplex settings are aggregated. If auto-negotiation is used, make sure it is resolved to the same speed for all ports in the port trunk.
  - To get a value of an Ethernet statistics counter for the port trunk, add the values of the counter of all ports in the port trunk.

4.1.2. Link Aggregation Limitations

- A port mirroring target port can not be member of a port trunk. However, a port mirroring source port can be member of a port trunk.
- A port working in QinQ mode cannot be a member of a port trunk.
- DHCP Relay Agent Client port cannot be a member of a port trunk.
- Load balancing between the links of a bundle is randomized and may not be ideal. For instance, if three 100Mbs links are aggregated, the resulting bandwidth of the port trunk may not be precisely 300Mbs.
- A Static MAC Address should not be configured to reside on an aggregated port – it may cause some frames destined for that address to be dropped.
- A secure port cannot be a member of a port trunk.

Note

The port trunk must be properly configured on both sides of the aggregated link. In switch-to-switch connections, if the configuration of both sides does not match (i.e. some ports are mistakenly not included in the port trunk), it will result in a loop. So the following procedure is strongly recommended to configure a port trunk.

1. *Disconnect or disable all the ports involved in the configuration, i.e. either being added to or removed from the port trunk.*
2. *Configure the port trunk on both switches.*
3. *Double-check the port trunk configuration on both switches.*
4. *Reconnect or re-enable the ports.*

If the port trunk is being configured while the ports are not disconnected or disabled, the port will be disabled for a few seconds automatically.

Note

The IEEE 802.3ad Link Aggregation standard requires all physical links in the port trunk to run at the same speed and in full-duplex mode. If this requirement is violated, the performance of the port trunk will drop.

The switch will raise an appropriate alarm, if such a speed/duplex mismatch is detected.

Note

RSTP dynamically calculates the path cost of the port trunk based on its aggregated bandwidth. However, if the aggregated ports are running at different speeds, the path cost may not be calculated correctly.

Note

Enabling RSTP is the best way for handling link redundancy in switch-to-switch connections composed of more than one physical link. If RSTP is enabled and increased bandwidth is not required, Link Aggregation should not be used because it may lead to a longer fail-over time.

4.2. Link Aggregation Configuration

The Link Aggregation menu is accessible from the main menu.

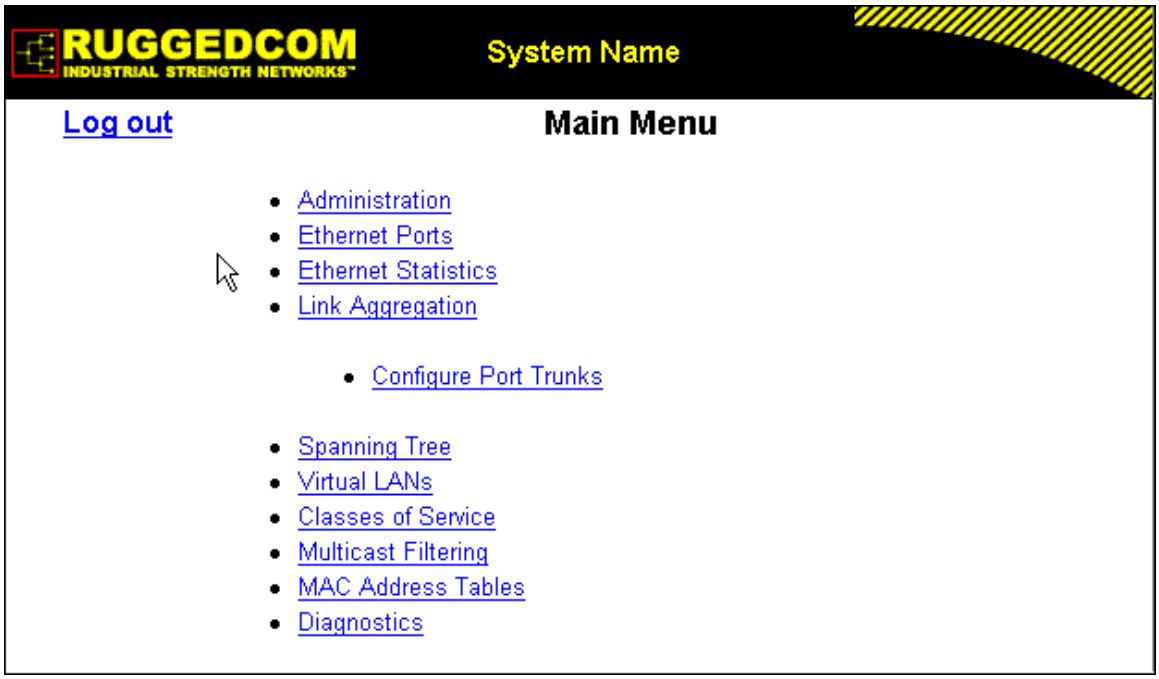

Figure 4.3. Link Aggregation Menu

4.2.1. Configuring Port Trunks

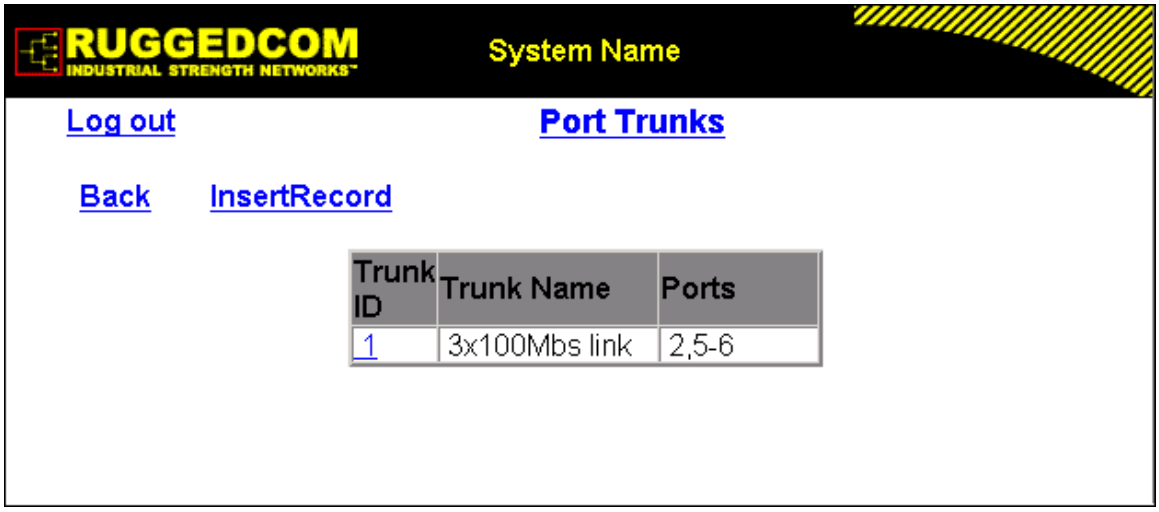

Figure 4.4. Port Trunk Table

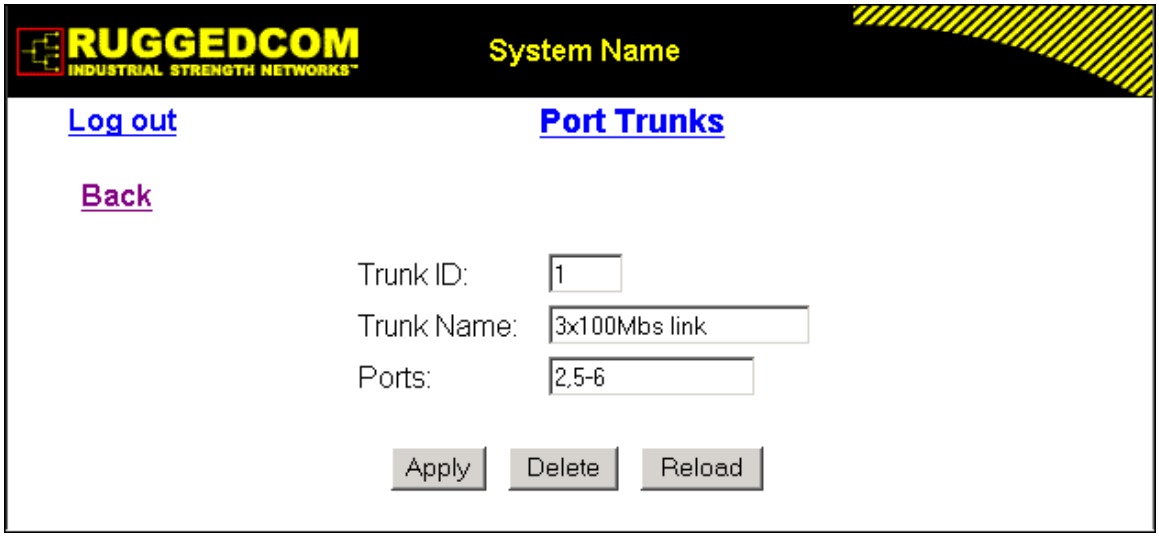

Figure 4.5. Port Trunk Form

Trunk ID

```
Synopsis: 1 to maximum number of port trunks
Default: 1
```

Trunk number. It doesn't affect port trunk operation in any way and is only used for identification.

Trunk Name

```
Synopsis: Any 19 characters
Default:
```

Provides a description of the aggregated link purpose.

Ports

```
Synopsis: Any combination of numbers valid for this parameter
Default: None
```

List of ports aggregated in the trunk.

5. Spanning Tree

The RuggedSwitch® family of Ethernet switches provides the latest in IEEE standard Spanning Tree functionality, including:

- Industry standard support of Rapid Spanning Tree (802.1D-2004), which features a compatibility mode with legacy STP (802.1D-1998)
- Industry standard support of Multiple Spanning Trees (802.1Q-2005), which is interoperable with both RSTP and legacy STP.
- RuggedCom RSTP feature enhancements (eRSTP™)
- Superior performance - RSTP will recognize a link failure and put an alternate port into forwarding within milliseconds.
- RSTP may be enabled on a per-port basis.
- Ports may be configured as edge ports, which allow rapid transitioning to the forwarding state for non-STP hosts.
- Path costs may be hard-configured or determined by port speed negotiation, in either the STP or RSTP style.
- Full bridge and port status displays provide a rich set of tools for performance monitoring and debugging.

Note

Historically, a device implementing STP on its ports has been referred to as a bridge. RuggedCom uses the terms "bridge" and "switch" synonymously.

- SNMP-manageable including newRoot and topologyChange traps.

5.1. RSTP Operation

The 802.1D Spanning Tree Protocol (STP) was developed to enable the construction of robust networks that incorporate redundancy while pruning the active topology of the network to prevent loops. While STP is effective, it requires that frame transfer halt after a link outage until all bridges in the network are guaranteed to be aware of the new topology. Using the values recommended by 802.1D, this period lasts 30 seconds.

The Rapid Spanning Tree Protocol (RSTP, IEEE 802.1w) was a further evolution of the 802.1D Spanning Tree Protocol. It replaced the settling period with an active handshake between bridges that guarantees the rapid propagation of topology information throughout the network. RSTP also offers a number of other significant innovations, including:

- Topology changes in RSTP can originate from and be acted upon by any designated bridges, leading to more rapid propagation of address information, unlike topology changes in STP, which must be passed to the root bridge before they can be propagated to the network.
- RSTP explicitly recognizes two blocking roles - Alternate and Backup Port - which are included in computations of when to learn and forward. STP, however, recognizes only one state - Blocking - for ports that should not forward.
- RSTP bridges generate their own configuration messages, even if they fail to receive any from the root bridge. This leads to quicker failure detection. STP, by contrast, must relay configuration

messages received on the root port out its designated ports. If an STP bridge fails to receive a message from its neighbor, it cannot be sure where along the path to the root a failure occurred.

- RSTP offers edge port recognition, allowing ports at the edge of the network to forward frames immediately after activation, while at the same time protecting them against loops.

While providing much better performance than STP, IEEE 802.1w RSTP still required up to several seconds to restore network connectivity when a topology change occurred.

A revised and highly optimized RSTP version was defined in the IEEE standard 802.1D-2004 edition. IEEE 802.1D-2004 RSTP reduces network recovery times to just milliseconds and optimizes RSTP operation for various scenarios.

ROS® supports IEEE 802.1D-2004 RSTP.

5.1.1. RSTP States and Roles

RSTP bridges have roles to play, either root or designated. One bridge - the Root Bridge - is the logical center of the network. All other bridges in the network are Designated bridges.

RSTP also assigns each port of the bridge a state and a role. The RSTP state describes what is happening at the port in relation to address learning and frame forwarding. The RSTP role basically describes whether the port is facing the center or the edges of the network and whether it can currently be used.

State

There are three RSTP states: Discarding, Learning and Forwarding.

The discarding state is entered when the port is first put into service. The port does not learn addresses in this state and does not participate in frame transfer. The port looks for RSTP traffic in order to determine its role in the network. When it is determined that the port will play an active part in the network, the state will change to learning.

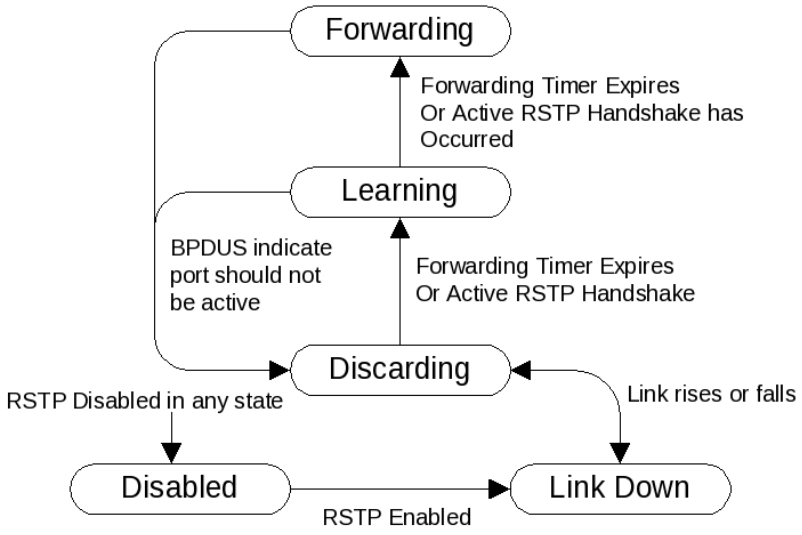

Figure 5.1. Bridge and Port States

The learning state is entered when the port is preparing to play an active part in the network. The port learns addresses in this state but does not participate in frame transfer. In a network of RSTP bridges, the time spent in this state is usually quite short. RSTP bridges operating in STP compatibility mode will spend six to 40 seconds in this state.

After “learning,” the bridge will place the port in the forwarding state. The port both learns addresses and participates in frame transfer while in this state.

Note

ROS introduces two more states - Disabled and Link Down. Introduced purely for purposes of management, these states may be considered subclasses of the RSTP Discarding state. The Disabled state refers to links for which RSTP has been disabled. The Link Down state refers to links for which RSTP is enabled but are currently down.

Role

There are four RSTP port roles: Root, Designated, Alternate and Backup.

If the bridge is not the root bridge, it must have a single Root Port. The Root Port is the “best” (i.e. quickest) way to send traffic to the root bridge.

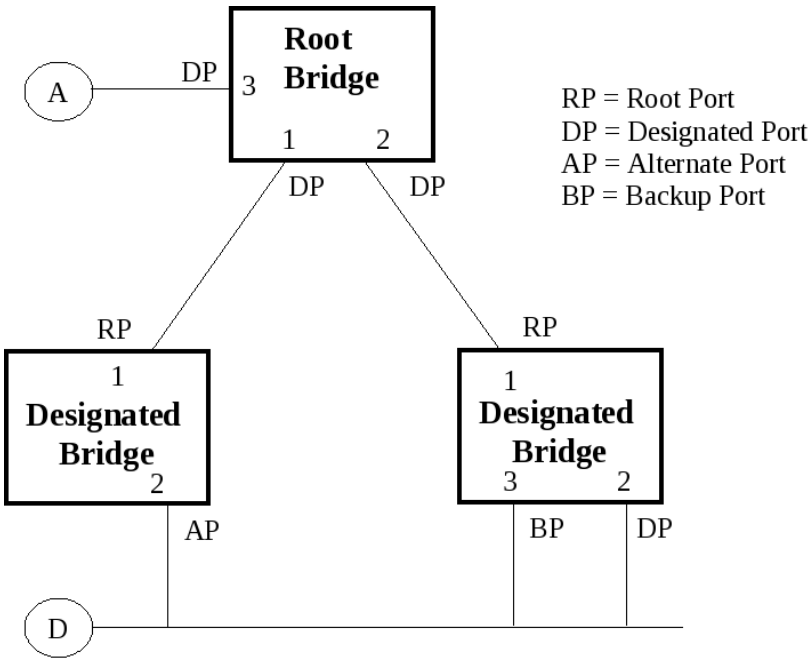

Figure 5.2. Bridge and Port Roles

A port is alternate when it receives a better message from another bridge on the LAN segment it is connected to. The message that an Alternate Port receives is better than the port itself would generate, but not good enough to convince it to become the Root Port. The port becomes the alternate to the current Root Port and will become the new Root Port should the current Root Port fail. The Alternate Port does not participate in the network.

A port is a Backup Port when it receives a better message from the LAN segment it is connected to, originating from another port on the same bridge. The port is a backup for another port on the bridge and will become active if that port fails. The Backup Port does not participate in the network.

5.1.2. Edge Ports

A port may be designated an Edge Port if it is directly connected to an end station. As such, it cannot create bridging loops in the network and can thus directly transition to forwarding, skipping the listening and learning stages.

Edge ports that receive configuration messages immediately lose their Edge Port status and become normal spanning tree ports. A loop created on an improperly connected edge port is thus quickly repaired.

Because an Edge Port services only end stations, topology change messages are not generated when its link toggles.

5.1.3. Point-to-Point and Multipoint Links

RSTP uses a peer-peer protocol called Proposing-Agreeing to ensure transitioning in the event of a link failure. This protocol is point-to-point and breaks down in multipoint situations, i.e. when more than two bridges operate on a shared media link.

If RSTP detects this circumstance (based upon the port's half duplex state after link up) it will switch off Proposing-Agreeing. The port must transition through the learning and forwarding states, spending one forward delay in each state.

There are circumstances in which RSTP will make an incorrect decision about the point-to-point state of the link simply by examining the half-duplex status, namely:

- The port attaches only to a single partner, but through a half-duplex link.
- The port attaches to a shared media hub through a full-duplex link. The shared media link attaches to more than one RSTP enabled bridge.

In such cases, the user may configure the bridge to override the half-duplex determination mechanism and force the link to be treated in the proper fashion.

5.1.4. Path and Port Costs

The STP path cost is the main metric by which root and designated ports are chosen[1]. The path cost for a designated bridge is the sum of the individual port costs of the links between the root bridge and that designated bridge. The port with the lowest path cost is the best route to the root bridge and is chosen as the root port.

How Port Costs Are Generated

Port costs can be generated either as a result of link auto-negotiation or manual configuration.

[1] In actuality the primary determinant for root port selection is the root bridge ID. Bridge ID is important mainly at network startup when the bridge with the lowest ID is elected as the root bridge. After startup (when all bridges agree on the root bridge's ID) the path cost is used to select root ports. If the path costs of candidates for the root port are the same, the ID of the peer bridge is used to select the port. Finally, if candidate root ports have the same path cost and peer bridge ID, the port ID of the peer bridge is used to select the root port. In all cases the lower ID, path cost or port ID is selected as the best.

When the link auto-negotiation method is used, the port cost is derived from the speed of the link. This method is useful when a well-connected network has been established. It can be used when the designer is not too concerned with the resultant topology as long as connectivity is assured.

Manual configuration is useful when the exact topology of the network must be predictable under all circumstances. The path cost can be used to establish the topology of the network exactly as the designer intends.

STP vs. RSTP Costs

The IEEE 802.1D-1998 specification limits port costs to values of 1 to 65536. It recommends that a path cost corresponding to the 1x109 / link speed be used. Designed at a time when 9600 bps links were state of the art, this method breaks down in modern use, as the method cannot represent a link speed higher than a gigabit per second.

In order to remedy this problem in future applications the IEEE 802.1w specification limits port costs to values of 1 to 200000, with a path cost corresponding to the 2x1012 / link speed.

RuggedCom bridges support interoperability with legacy STP bridges by selecting the style to use. In practice it makes no difference which style is used as long as it is applied consistently across the network, or if costs are manually assigned.

5.1.5. Bridge Diameter

The bridge diameter is the maximum number of bridges between any two possible points of attachment of end stations to the network.

The bridge diameter reflects the realization that topology information requires time to propagate hop by hop through a network. If configuration messages take too long to propagate end to end through the network, the result will be an unstable network.

There is a relationship between the bridge diameter and the maximum age parameter[2]. To achieve extended ring sizes, RuggedCom eRSTP™ uses an age increment of ¼ of a second. The value of the maximum bridge diameter is thus four times the configured maximum age parameter.

Note

Raise the value of the maximum age parameter if implementing very large bridged networks or rings.

5.2. MSTP Operation

The Multiple Spanning Tree (MST) algorithm and protocol provide greater control and flexibility than RSTP and legacy STP. MSTP (Multiple Spanning Tree Protocol) is an extension of RSTP, whereby multiple spanning trees may be maintained on the same bridged network. Data traffic

[2]The RSTP algorithm is as follows. STP configuration messages contain “age” information. Messages transmitted by the root bridge have an age of 0. As each subsequent designated bridge transmits the configuration message it must increase the age by at least 1 second. When the age exceeds the value of the maximum age parameter the next bridge to receive the message immediately discards it.

is allocated to one or another of several spanning trees by mapping one or more VLANs onto the network.

Note

The sophistication and utility of the Multiple Spanning Tree implementation on a given bridged network is proportional to the amount of planning and design invested in configuring MSTP.

If MSTP is activated on some or all of the bridges in a network with no additional configuration, the result will be a fully and simply connected network, but at best, the result will be the same as a network using only RSTP. Taking full advantage of the features offered by MSTP requires a potentially large number of configuration variables to be derived from an analysis of data traffic on the bridged network, and from requirements for load sharing, redundancy, and path optimization. Once these parameters have all been derived, it is also critical that they are consistently applied and managed across all bridges in an MST region.

5.2.1. MST Regions and Interoperability

In addition to supporting multiple spanning trees in a network of MSTP-capable bridges, MSTP is capable of interoperating with bridges that support only RSTP or legacy STP, without requiring any special configuration.

An MST region may be defined as the set of interconnected bridges whose MST Region Identification is identical (see section Section 5.4.3, “MST Region Identifier”). The interface between MSTP bridges and non-MSTP bridges, or between MSTP bridges with different MST Region Identification information, becomes part of an MST Region boundary.

Bridges outside an MST region will see the entire region as though it were a single (R)STP bridge; the internal detail of the MST region is hidden from the rest of the bridged network. In support of this, MSTP maintains separate ‘hop counters’ for spanning tree information exchanged at the MST region boundary versus that propagated inside the region. For information received at the MST region boundary, the (R)STP Message Age is incremented only once. Inside the region, a separate Remaining Hop Count is maintained, one for each spanning tree instance. The external Message Age parameter is referred to the (R)STP Maximum Age Time, whereas the internal Remaining Hop Counts are compared to an MST region-wide Maximum Hops parameter.

MSTI

An MSTI (Multiple Spanning Tree Instance) is one of sixteen independent spanning tree instances that may be defined in an MST region (not including the IST – see below). An MSTI is created by mapping a set of VLANs (in ROS, via the VLAN configuration) to a given MSTI ID. The same mapping must be configured on all bridges that are intended to be part of the MSTI. Moreover, all VLAN to MSTI mappings must be identical for all bridges in an MST region.

Note

ROS supports 16 MSTIs in addition to the IST

Each MSTI has a topology that is independent of every other. Data traffic originating from the same source and bound to the same destination but on different VLANs on different MSTIs may therefore travel a different path across the network.

IST

An MST region always defines an IST (Internal Spanning Tree). The IST spans the entire MST region, and carries all data traffic that is not specifically allocated (by VLAN) to a specific MSTI. The IST is always computed and is defined to be MSTI zero.

The IST is also the extension inside the MST region of the CIST (see below), which spans the entire bridged network, inside and outside of the MST region and all other RSTP and STP bridges, as well as any other MST regions.

CST

The CST (Common Spanning Tree) spans the entire bridged network, including MST regions and any connected STP or RSTP bridges. An MST region is seen by the CST as an individual bridge, with a single cost associated with its traversal.

CIST

The CIST (Common and Internal Spanning Tree) is the union of the CST and the ISTs in all MST regions. The CIST therefore spans the entire bridged network, reaching into each MST region via the latter's IST to reach every bridge on the network.

5.2.2. MSTP Bridge and Port Roles

5.2.2.1. Bridge Roles:

CIST Root

The CIST Root is the elected root bridge of the CIST (Common and Internal Spanning Tree), which spans all connected STP and RSTP bridges and MSTP regions.

CIST Regional Root

The root bridge of the IST within an MST region. The CIST Regional Root is the bridge within an MST region with the lowest cost path to the CIST Root. Note that the CIST Regional Root will be at the boundary of an MST region. Note also that it is possible for the CIST Regional Root to be the CIST Root.

MSTI Regional Root

The root bridge for an MSTI within an MST region. A root bridge is independently elected for each MSTI in an MST region.

5.2.2.2. Port Roles:

Each port on an MST bridge may have more than one role depending on the number and topology of spanning tree instances defined on the port.

CIST Port Roles

- The Root Port provides the minimum cost path from the bridge to the CIST Root via the CIST Regional Root. If the bridge itself happens to be the CIST Regional Root, the Root Port is also

the Master Port for all MSTIs (see below), and provides the minimum cost path to a CIST Root located outside the region.

- A Designated Port provides the minimum cost path from an attached LAN, via the bridge to the CIST Regional Root.
- Alternate and Backup Ports have the same sense that they do in RSTP, described in Section 5.1.1, “RSTP States and Roles”, under “Roles”, but relative to the CIST Regional Root.

MSTI Port Roles

For each MSTI on a bridge:

- The Root Port provides the minimum cost path from the bridge to the MSTI Regional Root, if the bridge itself is not the MSTI Regional Root.
- A Designated Port provides the minimum cost path from an attached LAN, via the bridge to the MSTI Regional Root.
- Alternate and Backup Ports have the same sense that they do in RSTP, described in Section 5.1.1, “RSTP States and Roles”, under “Roles”, but relative to the MSTI Regional Root.

The Master Port, which is unique in an MST region, is the CIST Root Port of the CIST Regional Root, and provides the minimum cost path to the CIST Root for all MSTIs.

Boundary Ports

A Boundary Port is a port on a bridge in an MST region that connects to either of: 1) a bridge belonging to a different MST region, or 2) a bridge supporting only RSTP or legacy STP. A Boundary Port blocks or forwards all VLANs from all MSTIs and the CIST alike. A Boundary Port may be:

- The CIST Root Port of the CIST Regional Root (and therefore also the MSTI Master Port).
- A CIST Designated Port, CIST Alternate / Backup Port, or Disabled. At the MST region boundary, the MSTI Port Role is the same as the CIST Port Role.

A Boundary Port connected to an STP bridge will send only STP BPDUs. One connected to an RSTP bridge need not refrain from sending MSTP BPDUs. This is made possible by the fact that the MSTP carries the CIST Regional Root Identifier in the field that RSTP parses as the Designated Bridge Identifier.

5.2.3. Benefits of MSTP

Despite the fact that MSTP is configured by default to arrive automatically at a spanning tree solution for each configured MSTI, advantages may be gained from influencing the topology of MSTIs in an MST region. The fact that the Bridge Priority and each port cost are configurable per MSTI (see sections Section 5.4.4, “Bridge MSTI Parameters” and Section 5.4.5, “Port MSTI Parameters”) makes it possible to control the topology of each MSTI within a region.

Load Balancing

MST can be used to balance data traffic load among (sets of) VLANs, enabling more complete utilization of a multiply interconnected bridged network.

A bridged network controlled by a single spanning tree will block redundant links by design, in order to avoid harmful loops. Using MSTP, however, any given link may have a different blocking state for each spanning tree instance (MSTI), as maintained by MSTP. Any given link, therefore, might be in blocking state for some VLANS and in forwarding state for other VLANs, depending on the mapping of VLANs to MSTIs.

It is possible to control the spanning tree solution for each MSTI, especially the set of active links for each tree, by manipulating, per MSTI, the bridge priority and the port costs of links in the network. If traffic is allocated judiciously to multiple VLANs, redundant interconnections in a bridged network which, using a single spanning tree, would have gone unused, can now be made to carry traffic.

Isolation of Spanning Tree Reconfiguration

A link failure in an MST region that does not affect the roles of Boundary ports will not cause the CST to be reconfigured, nor will the change affect other MST regions. This is due to the fact that MSTP information does not propagate past a region boundary.

MSTP versus PVST

An advantage of MSTP over the Cisco Systems Inc. proprietary PVST protocol is the ability to map multiple VLANs onto a single MSTI. Since each spanning tree requires processing and memory, the expense of keeping track of an increasing number of VLANs increases much more rapidly for PVST than for MSTP.

Compatibility with STP and RSTP

No special configuration is required for the bridges of an MST region to connect fully and simply to non-MST bridges on the same bridged network. Careful planning and configuration is, however, recommended in order to arrive at an optimal network.

5.2.4. Implementing MSTP on a Bridged Network

It is recommended that the configuration of MSTP on a network proceed in the sequence outlined below. Naturally, it is also recommended that network analysis and planning inform the steps of configuring the VLAN and MSTP parameters in particular.

Begin with a set of MSTP-capable Ethernet bridges, and MSTP disabled. For each bridge in the network:

1. Configure and enable RSTP (see sections Section 5.4.1, “Bridge RSTP Parameters” and Section 5.4.2, “Port RSTP Parameters”). Note that the Max Hops parameter in the Bridge RSTP Parameters menu is the maximum hop count for MSTP.
2. Create the VLANs that will be mapped to MSTIs (see the sections on VLAN Configuration).
3. Map VLANs to MSTIs (via the VLAN Configuration menus). Note that MSTP need not be enabled in order to map a VLAN to an MSTI. Note also that this mapping must be identical for each bridge that is to belong to the MST region.
4. Configure a Region Identifier and Revision Level. Note that these two items must be identical for each bridge in the MST region (see section Section 5.4.3, “MST Region Identifier”).

5. Verify that the Digest field in the MST Region Identifier menu is identical for each bridge in the MST region. If it is not, then the set of mappings from VLANs to MSTIs differs.
6. Configure Bridge Priority per MSTI (see section Section 5.4.4, “Bridge MSTI Parameters”).
7. Configure Port Cost and Priority per port and per MSTI (see section Section 5.4.5, “Port MSTI Parameters”).
8. Enable MSTP (see section Section 5.4.1, “Bridge RSTP Parameters”).

Note

Static VLANs must be used in an MSTP configuration. GVRP is not supported in this case.

5.3. RSTP Applications

5.3.1. RSTP in Structured Wiring Configurations

RSTP allows you to construct structured wiring systems in which connectivity is maintained in the event of link failures. For example, a single link failure of any of links A through N in Figure 5.3, “Example of a Structured Wiring Configuration”, would leave all the ports of bridges 555 through 888 connected to the network.

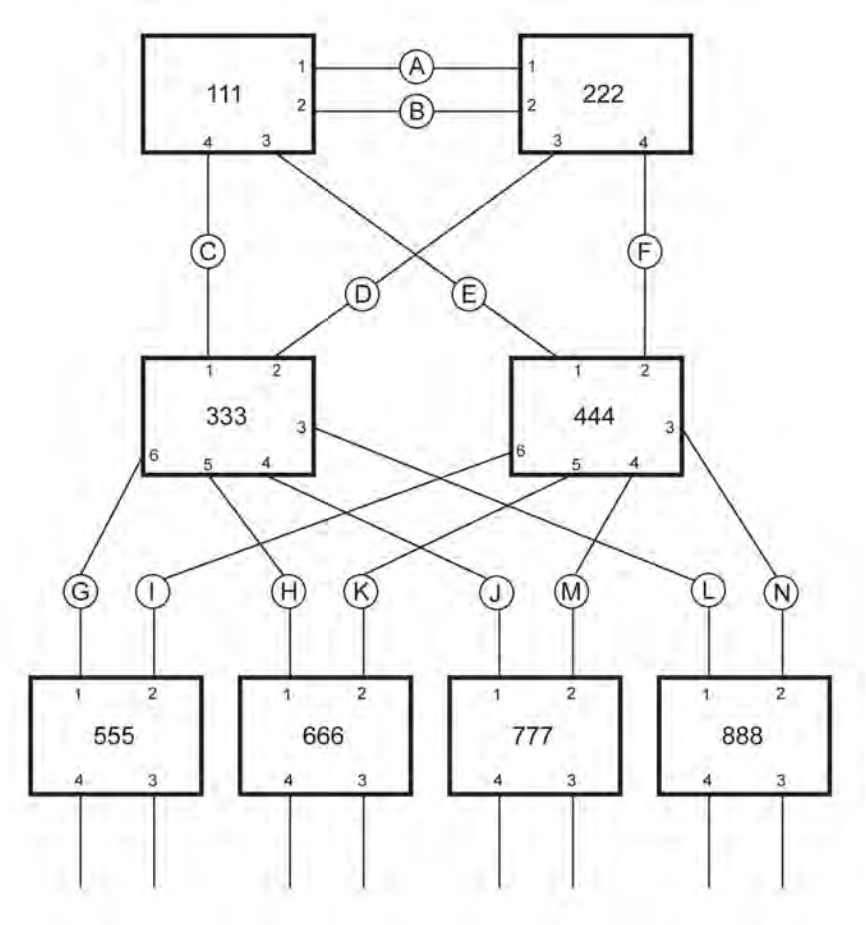

Figure 5.3. Example of a Structured Wiring Configuration

Procedure 5.1. Design Considerations for RSTP in Structured Wiring Configurations

1. **Select the design parameters for the network.**

   What are the requirements for robustness and network fail-over/recovery times? Are there special requirements for diverse routing to a central host computer? Are there any special port redundancy requirements?

2. **Identify required legacy support.**

   Are STP bridges used in the network? These bridges do not support rapid transitioning to forwarding. If these bridges are present, can they be re-deployed closer to the network edge?

3. **Identify edge ports and ports with half-duplex/shared media restrictions.**

   Ports that connect to host computers, IEDs and controllers may be set to edge ports in order to guarantee rapid transitioning to forwarding as well as to reduce the number of topology change notifications in the network. Ports with half-duplex/shared media restrictions require special attention in order to guarantee that they do not cause extended fail-over/recovery times.

4. **Choose the root bridge and backup root bridge carefully.**

   The root bridge should be selected to be at the concentration point of network traffic. Locate the backup root bridge adjacent to the root bridge. One strategy that may be used is to tune the bridge priority to establish the root bridge and then tune each bridge's priority to correspond to its distance from the root bridge.

5. **Identify desired steady state topology.**

   Identify the desired steady state topology taking into account link speeds, offered traffic and QOS. Examine of the effects of breaking selected links, taking into account network loading and the quality of alternate links.

6. **Decide upon port cost calculation strategy.**

   Select whether fixed or auto-negotiated costs should be used? Select whether the STP or RSTP cost style should be used.

7. **Calculate and configure priorities and costs.**

8. **Implement the network and test under load.**

5.3.2. RSTP in Ring Backbone Configurations

RSTP may be used in ring backbone configurations where rapid recovery from link failure is required. In normal operation, RSTP will block traffic on one of the links, for example as indicated by the double bars through link H in Figure 5.4, "Example of a Ring Backbone Configuration". In the event of a failure on link D, bridge 444 will unblock link H. Bridge 333 will communicate with the network through link F.

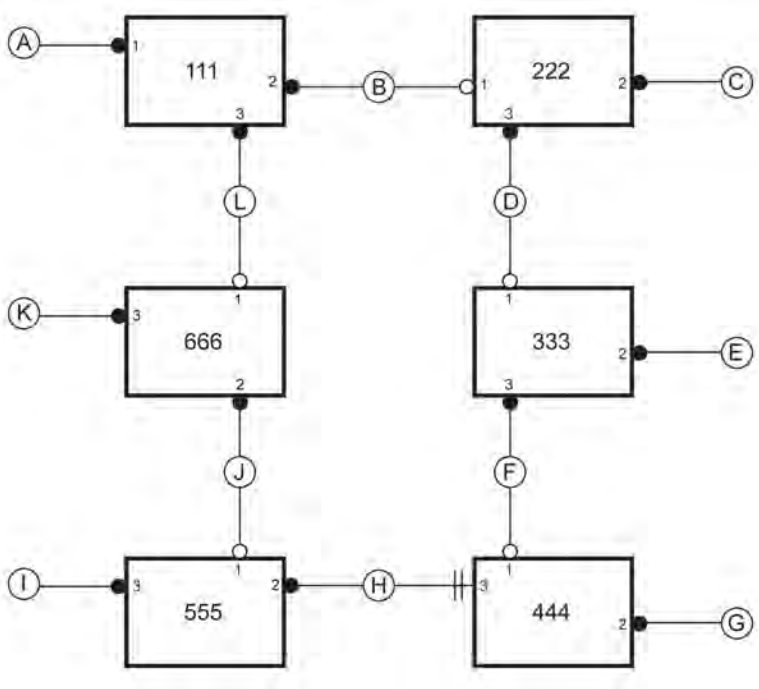

Figure 5.4. Example of a Ring Backbone Configuration

Procedure 5.2. Design Considerations for RSTP in Ring Backbone Configurations

1. **Select the design parameters for the network.**

   What are the requirements for robustness and network fail-over/recovery times? Typically, ring backbones are chosen to provide cost effective but robust network designs.

2. **Identify required legacy support and ports with half-duplex/shared media restrictions.**

   These bridges should not be used if network fail-over/recovery times are to be minimized.

3. **Identify edge ports**

   Ports that connect to host computers, IEDs and controllers may be set to edge ports in order to guarantee rapid transitioning to forwarding as well as to reduce the number of topology change notifications in the network.

4. **Choose the root bridge.**

   The root bridge can be selected to equalize either the number of bridges, number of stations or amount of traffic on either of its legs. It is important to realize that the ring will always be broken in one spot and that traffic always flows through the root.

5. **Assign bridge priorities to the ring.**

   The strategy that should be used is to assign each bridge's priority to correspond to its distance from the root bridge. If the root bridge is assigned the lowest priority of 0, the bridges on either side should use a priority of 4096 and the next bridges 8192 and so on. As there are 16 levels of bridge priority available, this method provides for up to 31 bridges in the ring.

6. **Implement the network and test under load.**

5.3.3. RSTP Port Redundancy

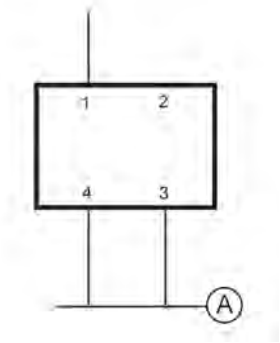

Figure 5.5. Port Redundancy

In cases where port redundancy is essential, RSTP allows more than one bridge port to service a LAN. For example, if port 3 is designated to carry the network traffic of LAN A, port 4 will block. Should an interface failure occur on port 3, port 4 would assume control of the LAN.

5.4. Spanning Tree Configuration

The Spanning Tree menu is accessible from the main menu.

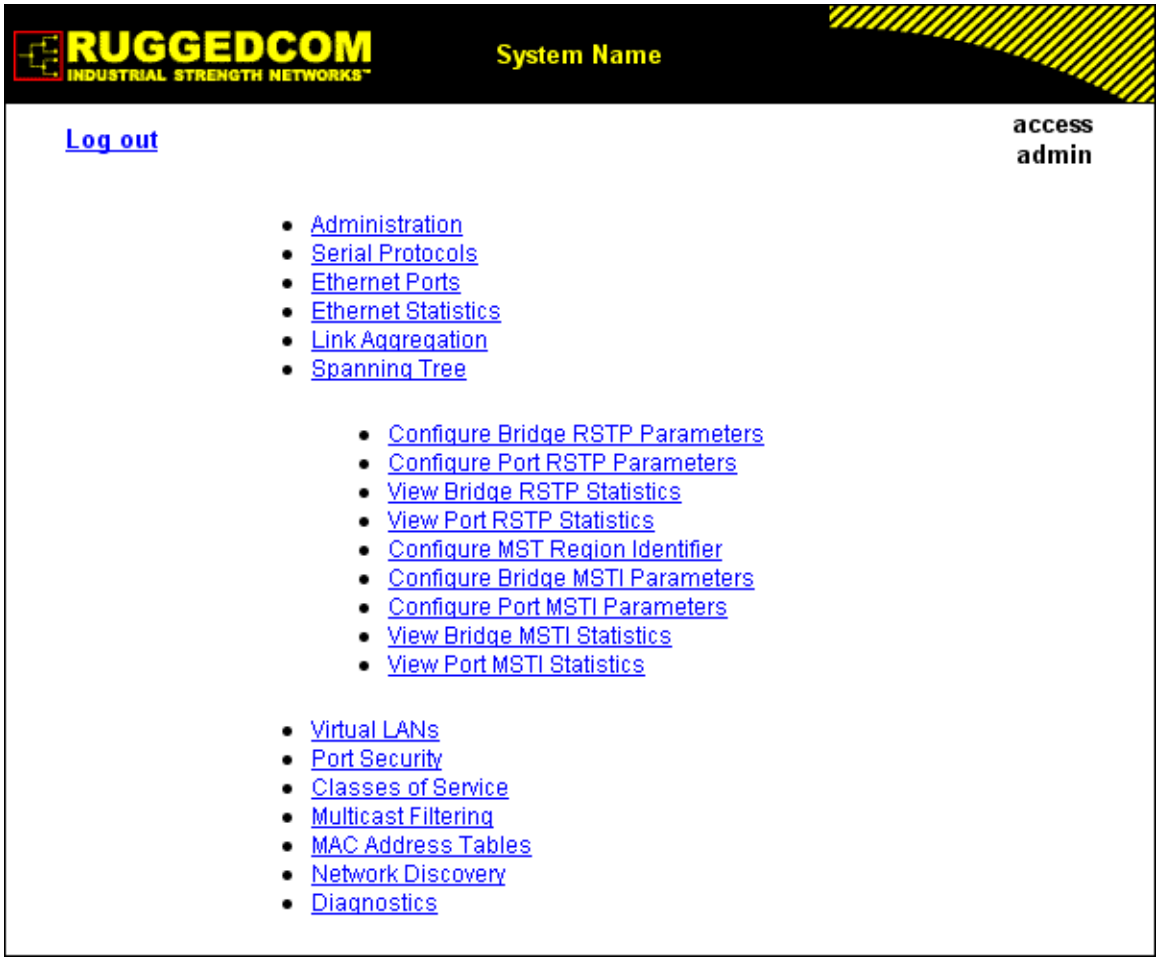

Figure 5.6. Spanning Tree Menu

5.4.1. Bridge RSTP Parameters

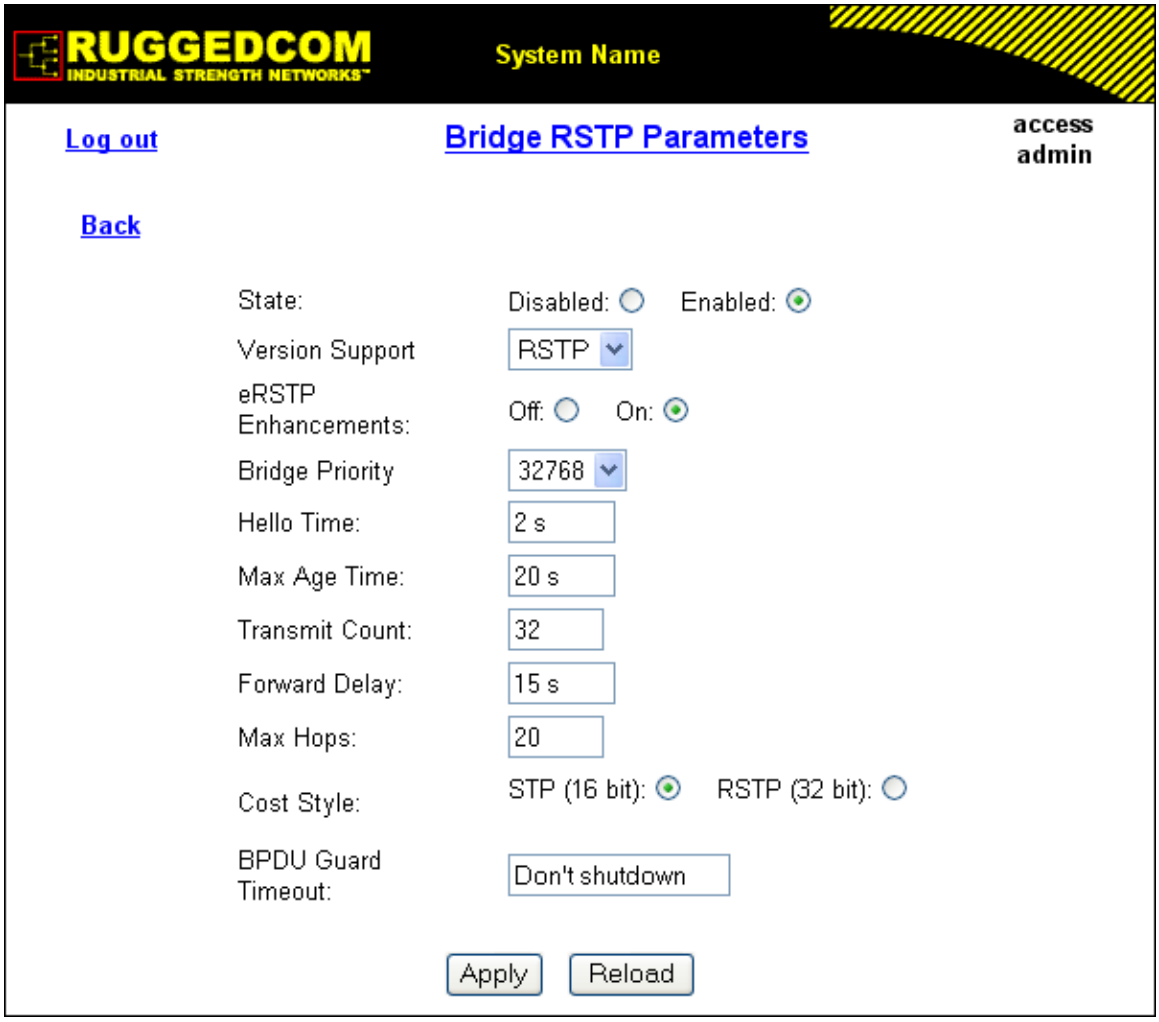

Figure 5.7. Bridge RSTP Parameters Form

State

```
Synopsis: { Disabled, Enabled }
Default: Enabled
```

Enable STP/RSTP/MSTP for the bridge globally. Note that for STP/RSTP/MSTP to be enabled on a particular port, it must be enabled both globally and per port.

Version Support

```
Synopsis: { STP, RSTP, MSTP }
Default: RSTP
```

Selects the version of Spanning Tree Protocol to support one of: STP, Rapid STP, or Multiple STP.

eRSTP Enhancements

```
Synopsis: { Off, On }
Default: On
```

Enable/disable RuggedCom proprietary eRSTP (enhanced RSTP) enhancements.

Bridge Priority

```
Synopsis: { 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768,
36864, 40960, 45056, 49152, 53248, 57344, 61440 }
Default: 32768
```

Bridge Priority provides a way to control the topology of the STP connected network. The desired Root and Designated bridges can be configured for a particular topology. The bridge with the lowest priority will become the root. In the event of a failure of the root bridge, the bridge with the next lowest priority will then become the root. Designated bridges that (for redundancy purposes) service a common LAN also use priority to determine which bridge is active. In this way, careful selection of Bridge Priorities can establish the path of traffic flows in normal and abnormal conditions.

Hello Time

```
Synopsis: 1 to 10
Default: 2 s
```

The time between configuration messages issued by the root bridge. Shorter hello times result in faster detection of topology changes at the expense of moderate increases in STP traffic.

Max Age Time

```
Synopsis: 6 to 40
Default: 20 s
```

The time for which a configuration message remains valid after being issued by the root bridge. Configure this parameter with care when many tiers of bridges exist, or when slow speed links (such as those used in WANs) are part of the network.

Transmit Count

```
Synopsis: 3 to 100
Default: 32
```

The maximum number of configuration messages on each port that may be sent in a special event (such as recovering from a failure or bringing up a new link). After the maximum number of messages is reached, RSTP will be limited to one message per second. Larger values allow the network to recover from failed links more quickly. If RSTP is being used in a ring architecture, the transmit count should be larger than the number of switches in the ring.

Forward Delay

```
Synopsis: 4 to 30
Default: 15 s
```

The amount of time a bridge spends learning MAC addresses on a rising port before beginning to forward traffic. Lower values allow the port to reach the forwarding state more quickly, but at the expense of flooding unlearned addresses to all ports.

Max Hops

```
Synopsis: 6 to 40
Default: 20
```

This parameter is only relevant for MSTP - ignore it otherwise.This parameter specifies the maximum possible bridge diameter inside an MST region. MSTP BPDUs propagating inside an MST region carry a time-to-live parameter that is decremented by every switch that propagates

the BPDU. If the maximum number of hops inside the region exceeds the configured maximum, BPDUs may be discarded due to their time-to-live information.

Cost Style

```
Synopsis: { STP (16 bit), RSTP (32 bit) }
Default: STP (16 bit)
```

This parameter selects the style of link costs to employ. STP uses 16-bit path costs based upon 1x10E9/link speed (4 for 1Gbps, 19 for 100 Mbps and 100 for 10 Mbps) whereas RSTP uses 32-bit costs based upon 2x10E13/link speed (20,000 for 1Gbps, 200,000 for 100 Mbps and 2,000,000 for 10 Mbps). Note that RSTP link costs are used only when the bridge version support is set to allow RSTP and the port does not migrate to STP.

BPDU Guard Timeout

```
Synopsis: 1 to 86400 s or { Until reset, Don't shutdown }
Default: Don't shutdown
```

The RSTP standard does not address network security. RSTP must process every received BPDU and take an appropriate action. This opens a way for an attacker to influence RSTP topology by injecting RSTP BPDUs into the network.

BPDU Guard is a feature that protects the network from BPDUs received by a port where RSTP capable devices are not expected to be attached. If a BPDU is received by a port for which 'Edge' parameter is set to 'TRUE' or RSTP is disabled, the port will be shutdown for the time period specified by this parameter.

DON'T SHUTDOWN - BPDU Guard is disabled

UNTIL RESET - port will remain shutdown until the port reset command is issued by the user

5.4.2. Port RSTP Parameters

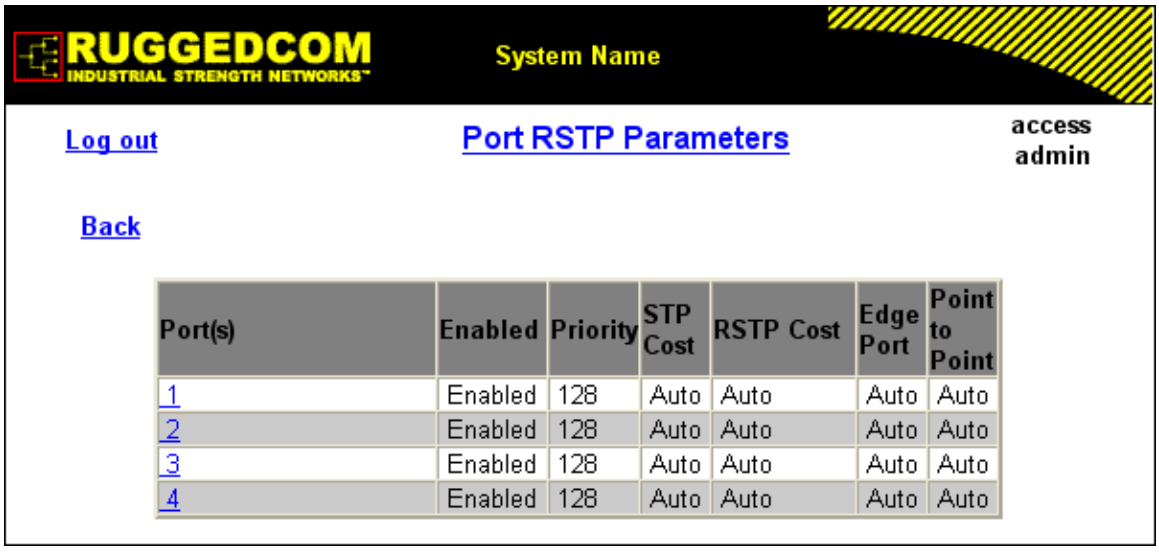

Port(s)	Enabled	Priority	STP Cost	RSTP Cost	Edge Port	Point to Point
1	Enabled	128	Auto	Auto	Auto	Auto
2	Enabled	128	Auto	Auto	Auto	Auto
3	Enabled	128	Auto	Auto	Auto	Auto
4	Enabled	128	Auto	Auto	Auto	Auto

Figure 5.8. Port RSTP Parameter Table

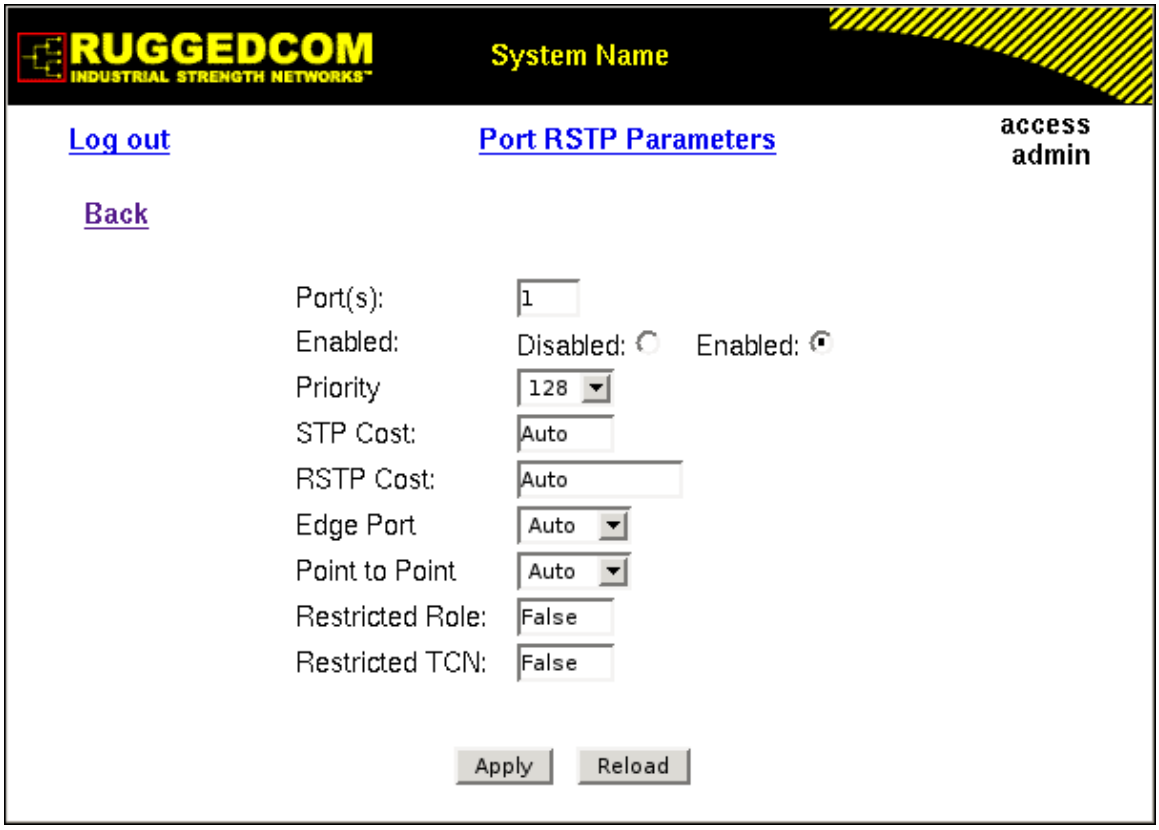

Figure 5.9. Port RSTP Parameter Form

Port(s)

```
Synopsis: Any combination of numbers valid for this parameter
```

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Enabled

```
Synopsis: { Disabled, Enabled }
Default: Enabled
```

Enabling STP activates the STP or RSTP protocol for this port per the configuration in the STP Configuration menu. STP may be disabled for the port ONLY if the port does not attach to an STP enabled bridge in any way. Failure to meet this requirement WILL result in an undetectable traffic loop in the network. A better alternative to disabling the port is to leave STP enabled but to configure the port as an edge port. A good candidate for disabling STP would be a port that services only a single host computer.

Priority

```
Synopsis: { 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 194,
208, 224, 240 }
Default: 128
```

Selects the STP port priority. Ports of the same cost that attach to a common LAN will select the port to be used based upon the port priority.

STP Cost

```
Synopsis: 0 to 65535 or { Auto }
```

```
Default: Auto
```

Selects the cost to use in cost calculations, when the Cost Style parameter is set to STP in the Bridge RSTP Parameters configuration. Setting the cost manually provides the ability to preferentially select specific ports to carry traffic over others. Leave this field set to "auto" to use the standard STP port costs as negotiated (4 for 1Gbps, 19 for 100 Mbps links and 100 for 10 Mbps links).

For MSTP, this parameter applies to both external and internal path cost.

RSTP Cost

```
Synopsis: 0 to 2147483647 or { Auto }
Default: Auto
```

Selects the cost to use in cost calculations, when the Cost Style parameter is set to RSTP in the Bridge RSTP Parameters configuration. Setting the cost manually provides the ability to preferentially select specific ports to carry traffic over others. Leave this field set to "auto" to use the standard RSTP port costs as negotiated (20,000 for 1Gbps, 200,000 for 100 Mbps links and 2,000,000 for 10 Mbps links).

For MSTP, this parameter applies to both external and internal path cost.

Edge Port

```
Synopsis: { False, True, Auto }
Default: Auto
```

Edge ports are ports that do not participate in the Spanning Tree, but still send configuration messages. Edge ports transition directly to frame forwarding without any listening and learning delays. The MAC tables of Edge ports do not need to be flushed when topology changes occur in the STP network. Unlike an STP disabled port, accidentally connecting an edge port to another port in the spanning tree will result in a detectable loop. The "Edgeness" of the port will be switched off and the standard RSTP rules will apply (until the next link outage).

Point to Point

```
Synopsis: { False, True, Auto }
Default: Auto
```

RSTP uses a peer-to-peer protocol that provides rapid transitioning on point-to-point links. This protocol is automatically turned off in situations where multiple STP bridges communicate over a shared (non point-to-point) LAN. The bridge will automatically take point-to-point to be true when the link is found to be operating in full-duplex mode. The point-to-point parameter allows this behavior or overrides it, forcing point-to-point to be true or false. Force the parameter true when the port operates a point-to-point link but cannot run the link in full-duplex mode. Force the parameter false when the port operates the link in full-duplex mode, but is still not point-to-point (e.g. a full-duplex link to an unmanaged bridge that concentrates two other STP bridges).

Restricted Role

```
Synopsis: { True or False }
Default: False
```

A boolean value set by management. If TRUE, causes the Port not to be selected as the Root Port for the CIST or any MSTI, even if it has the best spanning tree priority vector. Such a Port will be selected as an Alternate Port after the Root Port has been selected. This parameter should be FALSE by default. If set, it can cause a lack of spanning tree connectivity. It is set by a network administrator to prevent bridges that are external to a core region of the network from influencing the spanning tree active topology. This may be necessary, for example, if those bridges are not under the full control of the administrator.

Restricted TCN

```
Synopsis: { True or False }
Default: False
```

A boolean value set by management. If TRUE, it causes the Port not to propagate received topology change notifications and topology changes to other Ports. If set, it can cause temporary loss of connectivity after changes in a spanning tree's active topology as a result of persistent, incorrectly learned, station location information. It is set by a network administrator to prevent bridges that are external to a core region of the network from causing address flushing in that region. This may be necessary, for example, if those bridges are not under the full control of the administrator or if the MAC_Operational status parameter for the attached LANs transitions frequently.

5.4.3. MST Region Identifier

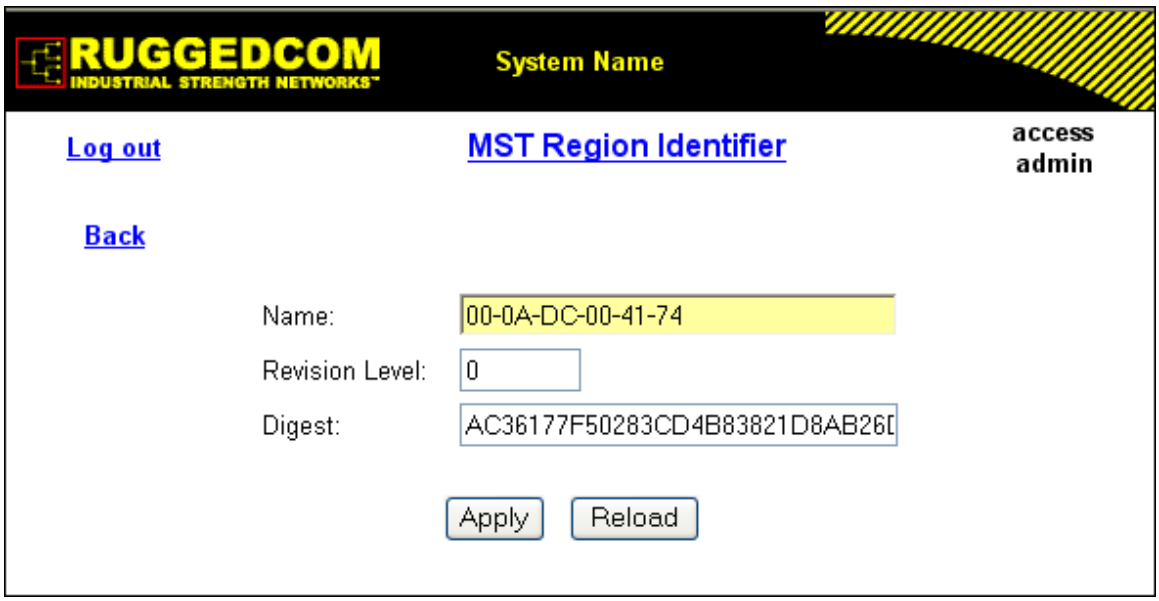

Figure 5.10. MST Region Identifier Form

Name

```
Synopsis: Any 32 characters
Default: 00-0A-DC-00-41-74
```

Variable length text string. You must configure an identical region name on all switches you want to be in the same MST region.

Revision Level

```
Synopsis: 0 to 65535
Default: 0
```

Use this parameter, if you want to create a new region from a subset of switches in a current region, while maintaining the same region name.

Digest

```
Synopsis: 32 hex characters
```

This is a read-only parameter and should be only used for network troubleshooting. In order to ensure consistent VLAN-to-instance mapping, it is necessary for the protocol to be able to

exactly identify the boundaries of the MST regions. For that purpose, the characteristics of the region are included in BPDUs. There is no need to propagate the exact VLAN-to-instance mapping in the BPDUs because switches only need to know whether they are in the same region as a neighbor. Therefore, only this 16-octet digest created from the VLAN-to-instance mapping is sent in BPDUs.

5.4.4. Bridge MSTI Parameters

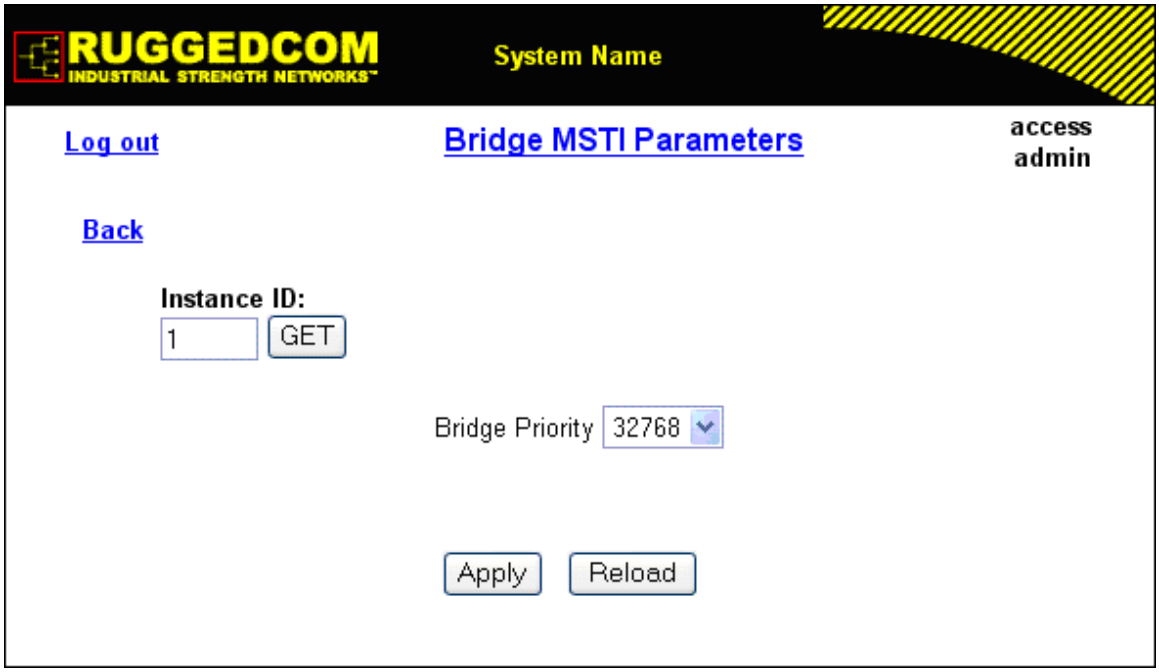

Figure 5.11. Bridge MSTI Parameters

Instance ID

```
Synopsis: 0 to 16
Default: 1
```

The Instance ID refers to the MSTI (Multiple Spanning Tree Instance) ID. Specify an Instance ID and select GET in order to load the parameters of the page corresponding to the selected MSTI. Changes to parameters that are subsequently applied will apply to the selected Instance ID. Note: Bridge Parameters for the IST (MSTI zero) are accessible via the Bridge RSTP Parameters menu (see section Section 5.4.1, “Bridge RSTP Parameters”).

Bridge Priority

```
Synopsis: { 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768,
36864, 40960, 45056, 49152, 53248, 57344, 61440 }
Default: 32768
```

Bridge Priority provides a way to control the topology of the STP connected network. The desired Root and Designated bridges can be configured for a particular topology. The bridge with the lowest priority will become root. In the event of a failure of the root bridge, the bridge with the next lowest priority will then become root. Designated bridges that (for redundancy purposes) service a common LAN also use priority to determine which bridge is active. In this way careful selection of Bridge Priorities can establish the path of traffic flows in both normal and abnormal conditions.

5.4.5. Port MSTI Parameters

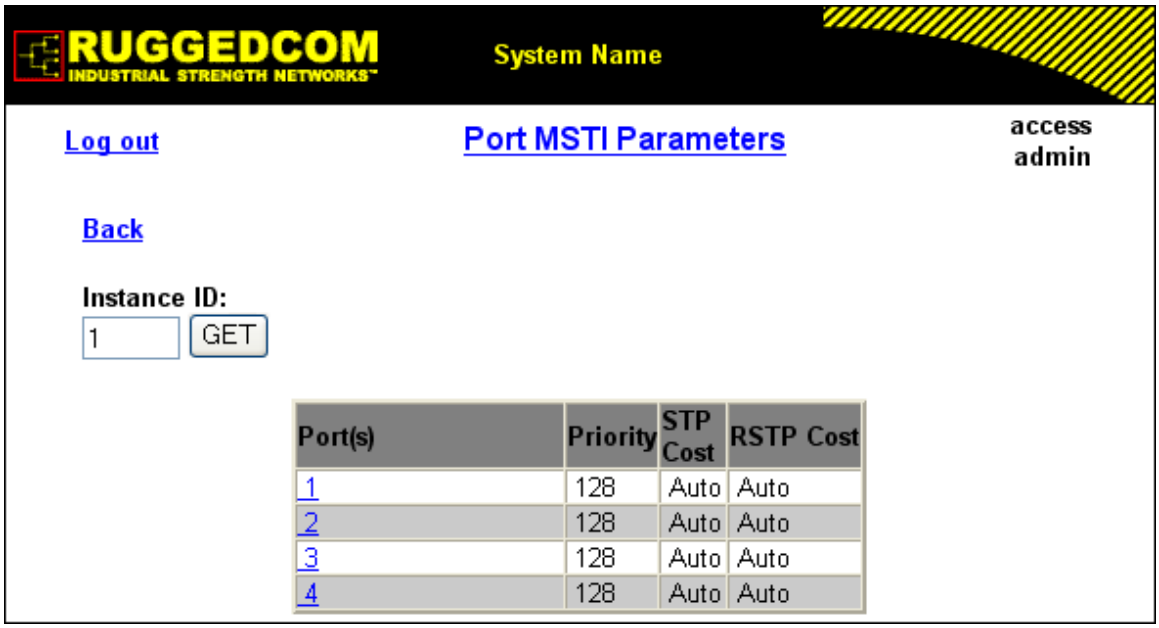

Figure 5.12. Port MSTI Parameter Table

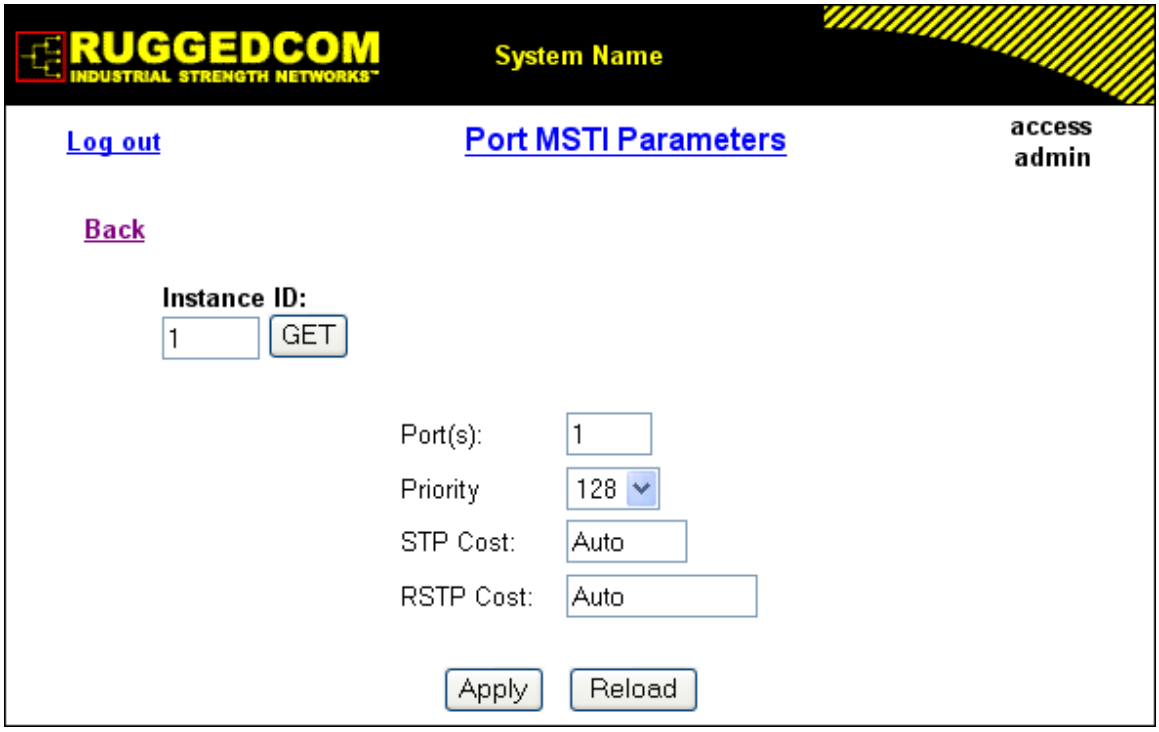

Figure 5.13. Port MSTI Parameter Form

Instance ID

```
Synopsis: 0 to 16
Default: 1
```

The Instance ID refers to the MSTI (Multiple Spanning Tree Instance) ID. Specify an Instance ID and select GET in order to load parameters corresponding to the selected MSTI. Changes to parameters that are subsequently applied will apply to the selected Instance ID. Note: Port Parameters for the IST (MSTI zero), are accessible via the Port RSTP Parameters menu (see section Section 5.4.2, “Port RSTP Parameters”).

Port(s)

```
Synopsis: Any combination of numbers valid for this parameter
```

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Priority

```
Synopsis: { 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 194,
208, 224, 240 }
Default: 128
```

Selects the STP port priority. Ports of the same cost that attach to a common LAN will select the port to be used based on the port priority.

STP Cost

```
Synopsis: 0 to 65535 or { Auto }
Default: Auto
```

Selects the cost to use in cost calculations when the Cost Style parameter is set to STP in the Bridge RSTP Parameters configuration. Setting the cost manually provides the ability to preferentially select specific ports to carry traffic over others. Leave this field set to "auto" to use the standard STP port costs as negotiated (4 for 1Gbps, 19 for 100 Mbps links and 100 for 10 Mbps links). For MSTP, this parameter applies to both external and internal path cost.

RSTP Cost

```
Synopsis: 0 to 2147483647 or { Auto }
Default: Auto
```

Selects the cost to use in cost calculations when the Cost Style parameter is set to RSTP in the Bridge RSTP Parameters configuration. Setting the cost manually provides the ability to preferentially select specific ports to carry traffic over others. Leave this field set to "auto" to use the standard RSTP port costs as negotiated (20,000 for 1Gbps, 200,000 for 100 Mbps links and 2,000,000 for 10 Mbps links). For MSTP, this parameter applies to both external and internal path cost.

5.5. Spanning Tree Statistics

5.5.1. Bridge RSTP Statistics

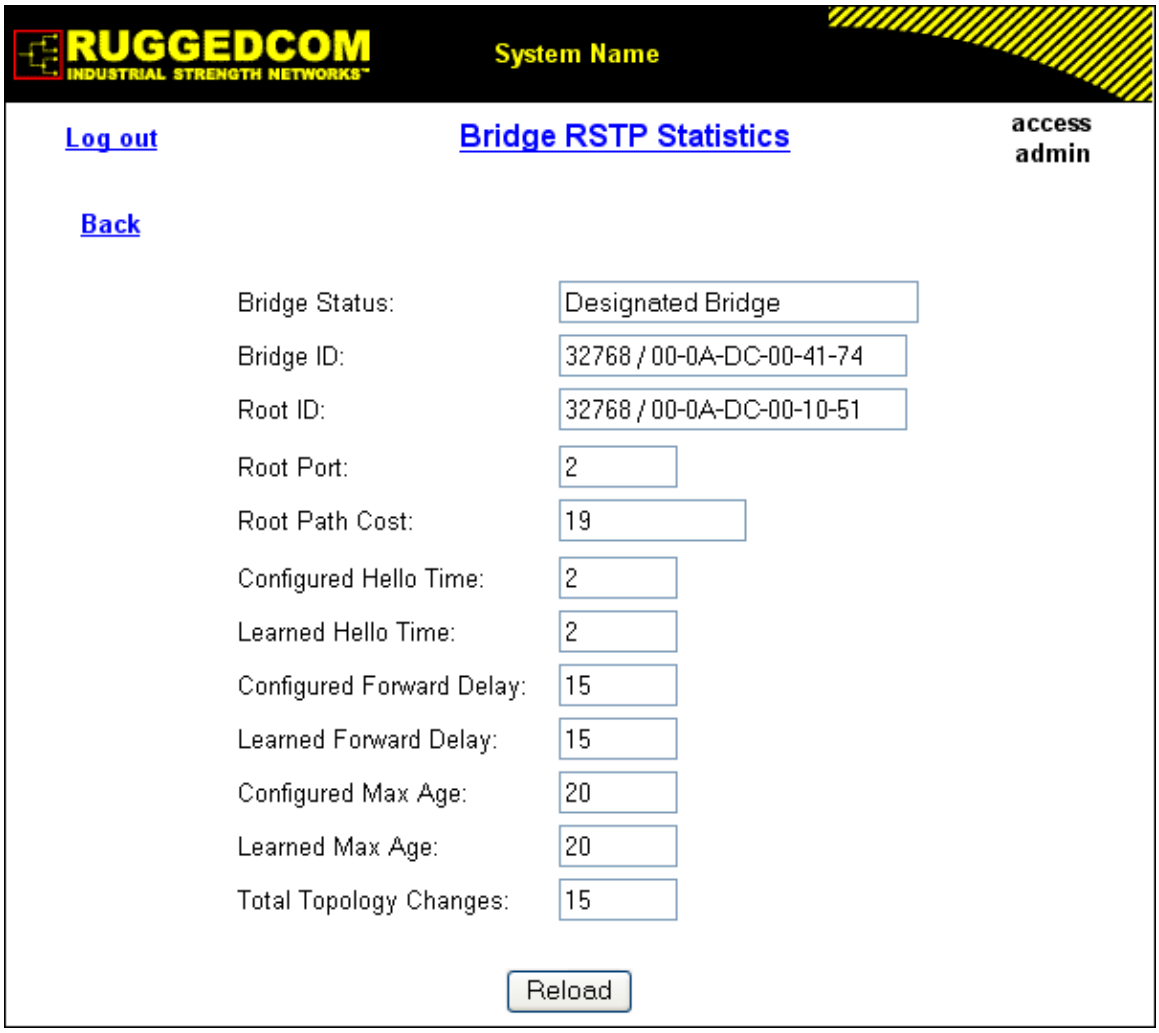

Figure 5.14. Bridge RSTP Statistics Form

Bridge Status

`Synopsis: { <empty string>, Designated Bridge, Not Designated For Any LAN, Root Bridge }`

Spanning Tree status of the bridge. The status may be root or designated. This field may display "Not designated For Any LAN" if the bridge is not the designated bridge for any of its ports.

Bridge ID

`Synopsis: $$ / ##-##-##-##-##-## where $$ is 0 to 65535, ## is 0 to FF`

Bridge Identifier of this bridge.

Root ID

`Synopsis: $$ / ##-##-##-##-##-## where $$ is 0 to 65535, ## is 0 to FF`

Bridge Identifier of the root bridge.

Root Port

`Synopsis: 0 to 65535 or { <empty string>}`

If the bridge is designated, this is the port that provides connectivity towards the root bridge of the network.

Root Path Cost

`Synopsis: 0 to 4294967295`

The total cost of the path to the root bridge, composed of the sum of the costs of each link in the path. If custom costs have not been configured. 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports will contribute 100. For the CIST instance of MSTP, this is an external root path cost, which is the cost of the path from the IST root (i.e. regional root) bridge to the CST root (i.e. network "global" root) bridge.

Configured Hello Time

`Synopsis: 0 to 65535`

The configured Hello time from the Bridge RSTP Parameters menu.

Learned Hello Time

`Synopsis: 0 to 65535`

The actual Hello time provided by the root bridge as learned in configuration messages. This time is used in designated bridges.

Configured Forward Delay

`Synopsis: 0 to 65535`

The configured Forward Delay time from the Bridge RSTP Parameters menu.

Learned Forward Delay

`Synopsis: 0 to 65535`

The actual Forward Delay time provided by the root bridge as learned in configuration messages. This time is used in designated bridges.

Configured Max Age

`Synopsis: 0 to 65535`

The configured Maximum Age time from the Bridge RSTP Parameters menu.

Learned Max Age

`Synopsis: 0 to 65535`

The actual Maximum Age time provided by the root bridge as learned in configuration messages. This time is used in designated bridges.

Total Topology Changes

`Synopsis: 0 to 65535`

A count of topology changes in the network, as detected on this bridge through link failures or as signaled from other bridges. Excessively high or rapidly increasing counts signal network problems.

5.5.2. Port RSTP Statistics

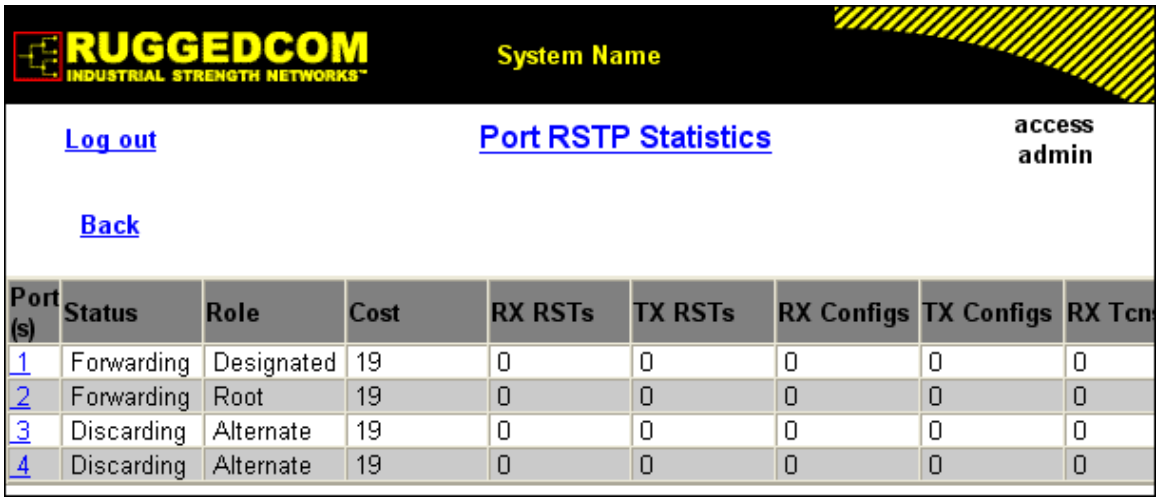

Port (s)	Status	Role	Cost	RX RSTs	TX RSTs	RX Configs	TX Configs	RX Tcns
1	Forwarding	Designated	19	0	0	0	0	0
2	Forwarding	Root	19	0	0	0	0	0
3	Discarding	Alternate	19	0	0	0	0	0
4	Discarding	Alternate	19	0	0	0	0	0

Figure 5.15. Port RSTP Statistics Table

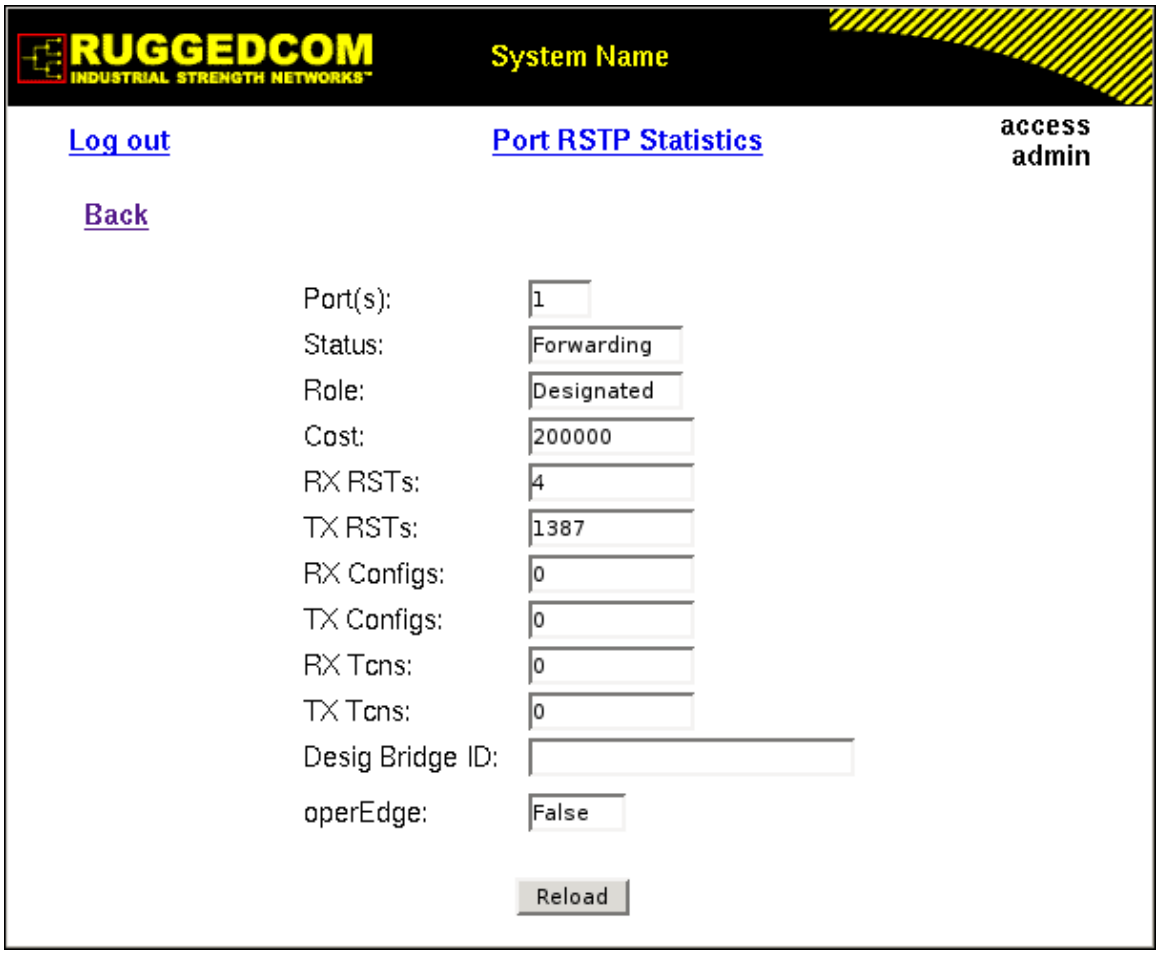

Figure 5.16. Port RSTP Statistics Form

Port(s)

Synopsis: Any combination of numbers valid for this parameter

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Status

`Synopsis: { Disabled, Listening, Learning, Forwarding, Blocking, Link Down, Discarding }`

The status of this port in the Spanning Tree. This may be one of the following:

Disabled - STP is disabled on this port.

Link Down - STP is enabled on this port but the link is down.

Discarding - The link is not used in the STP topology but is standing by.

Learning - The port is learning MAC addresses in order to prevent flooding when it begins forwarding traffic.

Forwarding - The port is forwarding traffic.

Role

`Synopsis: { <empty string>, Root, Designated, Alternate, Backup, Master }`

The role of this port in the Spanning Tree. This may be one of the following:

Designated - The port is designated for (i.e. carries traffic towards the root for) the LAN it is connected to.

Root - The single port on the bridge, which provides connectivity towards the root bridge.

Backup - The port is attached to a LAN that is serviced by another port on the bridge. It is not used but is standing by.

Alternate - The port is attached to a bridge that provides connectivity to the root bridge. It is not used but is standing by.

Cost

`Synopsis: 0 to 4294967295`

Cost offered by this port. If the Bridge RSTP Parameters Cost Style is set to STP, 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports contribute 100. If the Cost Style is set to RSTP, 1Gbps will contribute 20,000, 100 Mbps ports will contribute a cost of 200,000 and 10 Mbps ports contribute a cost of 2,000,000. Note that even if the Cost Style is set to RSTP, a port that migrates to STP will have its cost limited to a maximum of 65535.

RX RSTs

`Synopsis: 0 to 4294967295`

The count of RSTP configuration messages received on this port.

TX RSTs

`Synopsis: 0 to 4294967295`

The count of RSTP configuration messages transmitted on this port.

RX Configs

`Synopsis: 0 to 4294967295`

The count of STP configuration messages received on this port.

TX Configs

`Synopsis: 0 to 4294967295`

The count of STP configuration messages transmitted on this port.

RX Tcns

`Synopsis: 0 to 4294967295`

The count of configuration change notification messages received on this port. Excessively high or rapidly increasing counts signal network problems.

TX Tcns

`Synopsis: 0 to 4294967295`

The count of configuration messages transmitted from this port.

Desig Bridge ID

`Synopsis: $$ / ##-##-##-##-##-## where $$ is 0 to 65535, ## is 0 to FF`

Provided on the root ports of designated bridges, the Bridge Identifier of the bridge this port is connected to.

operEdge

`Synopsis: { True or False }`

Whether or not the port is operating as an edge port.

5.5.3. Bridge MSTI Statistics

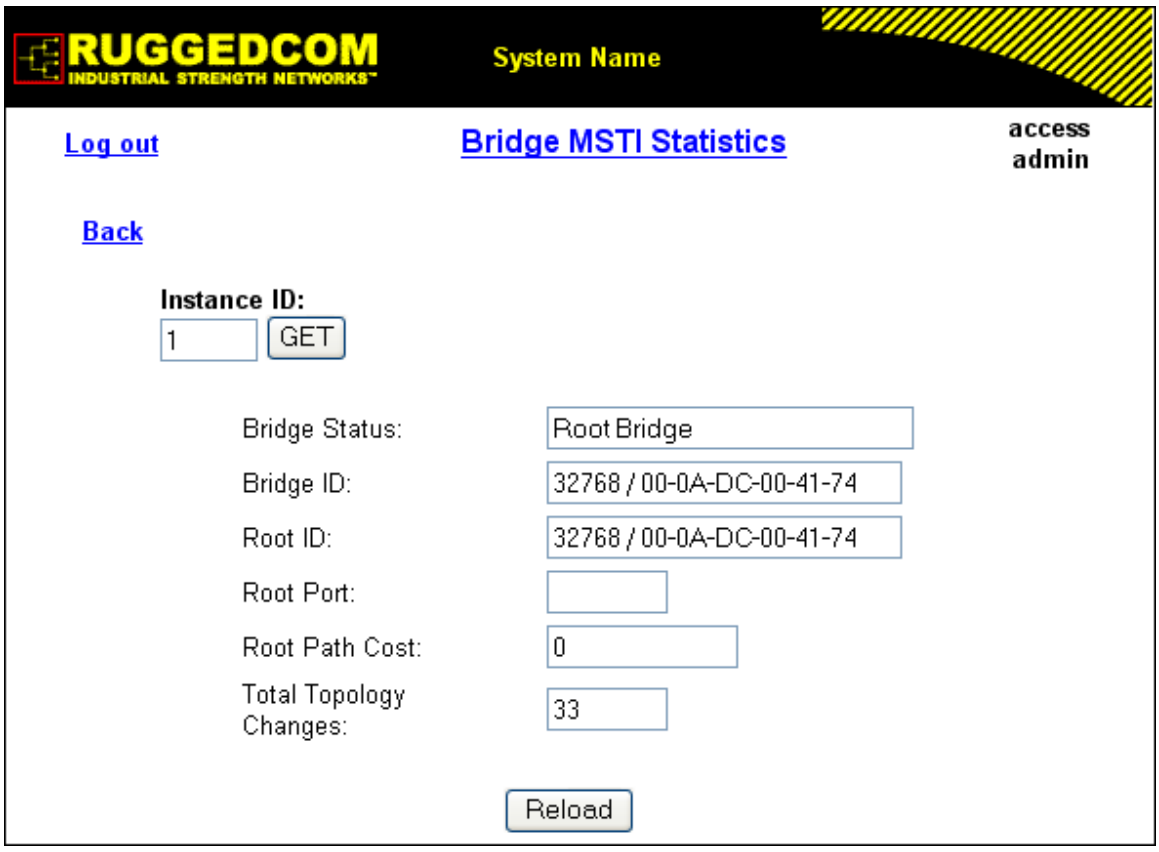

Figure 5.17. Bridge MSTI Statistics Form

Instance ID

`Synopsis: 0 to 16`

```
Default: 1
```

The Instance ID refers to the MSTI (Multiple Spanning Tree Instance) ID. Specify an Instance ID and select GET in order to load parameters corresponding to the selected MSTI. Note: Bridge Statistics for the IST (MSTI zero), are accessible via the Bridge RSTP Statistics menu (see section Section 5.5.1, “Bridge RSTP Statistics”).

Bridge Status

```
Synopsis: { <empty string>, Designated Bridge, Not Designated For Any
LAN, Root Bridge }
```

Spanning Tree status of the bridge. The status may be root or designated. This field may display "Not designated For Any LAN" if the bridge is not the designated bridge for any of its ports.

Bridge ID

```
Synopsis: $$ / ##-##-##-##-##-## where $$ is 0 to 65535, ## is 0 to FF
```

Bridge Identifier of this bridge.

Root ID

```
Synopsis: $$ / ##-##-##-##-##-## where $$ is 0 to 65535, ## is 0 to FF
```

Bridge Identifier of the root bridge.

Root Port

```
Synopsis: 0 to 65535 or { <empty string>}
```

If the bridge is designated, this is the port that provides connectivity towards the root bridge of the network.

Root Path Cost

```
Synopsis: 0 to 4294967295
```

The total cost of the path to the root bridge composed of the sum of the costs of each link in the path. If custom costs have not been configured. 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports will contribute 100 to this figure. For the CIST instance of MSTP, this is an external root path cost, which is the cost of the path from the IST root (i.e. regional root) bridge to the CST root (i.e. network "global" root) bridge.

Total Topology Changes

```
Synopsis: 0 to 65535
```

A count of topology changes in the network, as detected on this bridge through link failures or as signaled from other bridges. Excessively high or rapidly increasing counts signal network problems.

5.5.4. Port MSTI Statistics

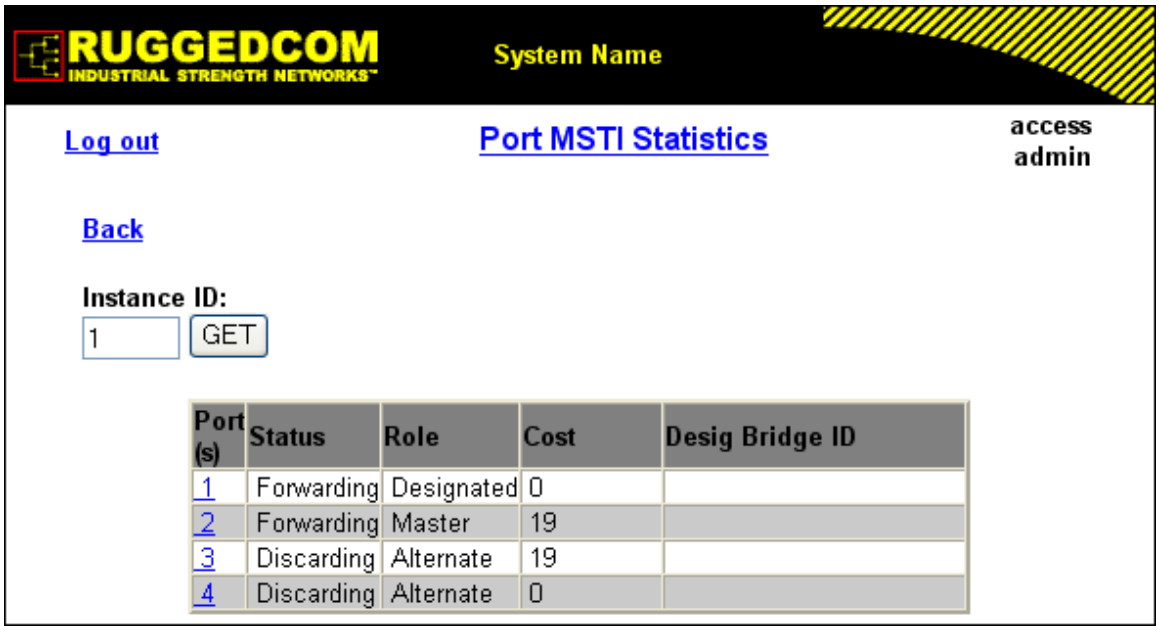

Figure 5.18. Port MSTI Statistics Table

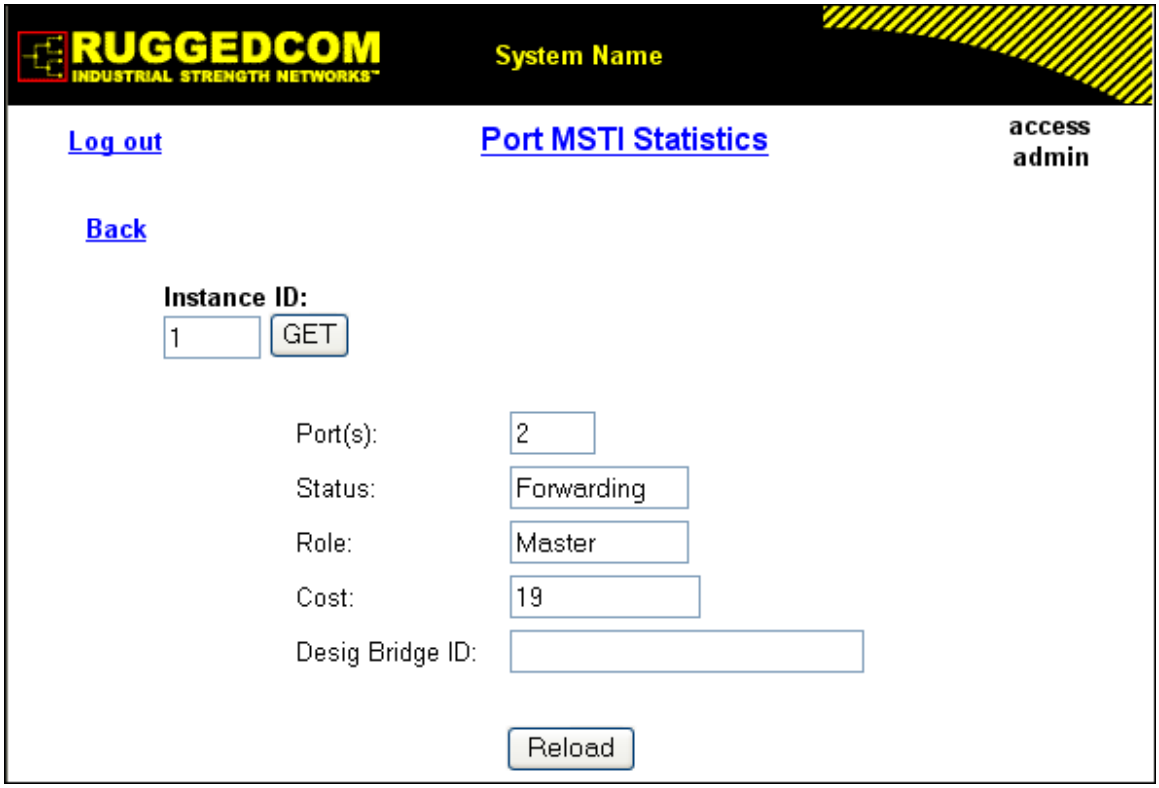

Figure 5.19. Port MSTI Statistics Form

Instance ID

```
Synopsis: 1 to 16
Default: 1
```

The Instance ID refers to the MSTI (Multiple Spanning Tree Instance) ID. Specify an Instance ID and select GET in order to load parameters corresponding to the selected MSTI. Note: Port Statistics for the IST (MSTI zero), are accessible via the Port RSTP Statistics menu (see section Section 5.5.2, “Port RSTP Statistics”).

Port(s)

`Synopsis: Any combination of numbers valid for this parameter`

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Status

`Synopsis: { Disabled, Listening, Learning, Forwarding, Blocking, Link Down, Discarding }`

The status of this port in the Spanning Tree. This may be one of the following:

Disabled - STP is disabled on this port.

Link Down - STP is enabled on this port but the link is down.

Discarding - The link is not used in the STP topology but is standing by.

Learning - The port is learning MAC addresses in order to prevent flooding when it begins forwarding traffic.

Forwarding - The port is forwarding traffic.

Role

`Synopsis: { <empty string>, Root, Designated, Alternate, Backup, Master }`

The role of this port in the Spanning Tree. This may be one of the following:

Designated - The port is designated for (i.e. carries traffic towards the root for) the LAN it is connected to.

Root - The single port on the bridge, which provides connectivity towards the root bridge.

Backup - The port is attached to a LAN that is serviced by another port on the bridge. It is not used but is standing by.

Alternate - The port is attached to a bridge that provides connectivity to the root bridge. It is not used but is standing by.

Master - Only exists in MSTP. The port is an MST region boundary port and the single port on the bridge, which provides connectivity for the Multiple Spanning Tree Instance towards the Common Spanning Tree root bridge (i.e. this port is the root port for the Common Spanning Tree Instance).

Cost

`Synopsis: 0 to 4294967295`

Cost offered by this port. If the Bridge RSTP Parameters Cost Style is set to STP, 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports contribute. If the Cost Style is set to RSTP, 1Gbps will contribute 20,000, 100 Mbps ports will contribute a cost of 200,000 and 10 Mbps ports contribute a cost of 2,000,000. Note that even if the Cost Style is set to RSTP, a port that migrates to STP will have its cost limited to a maximum of 65535.

Desig Bridge ID

`Synopsis: $$ / ##-##-##-##-##-## where $$ is 0 to 65535, ## is 0 to FF`

Provided on the root ports of designated bridges, the Bridge Identifier of the bridge this port is connected to.

5.6. Troubleshooting

Problem One

When I connect a new port, the network locks up. The port status LEDs are flashing madly.

Occasionally, the network seems to experience a lot of flooding. All the ports seem to experience significant traffic. The problem lasts a few seconds and then goes away.

One of my switches displays a strange behavior where the root port hops back and forth between two switch ports and never settles down.

Is it possible that one of the switches in the network or one of the ports on a switch in the network has STP disabled and accidentally connects to another switch? If this has occurred, then a traffic loop has been formed.

If the problem appears to be transient in nature, it is possible that ports that are part of the spanning tree have been configured as edge ports. After the link layers have come up on edge ports, STP will directly transition them (perhaps improperly) to the forwarding state. If an RSTP configuration message is then received, the port will be returned to blocking. A traffic loop may be formed for the length of time the port was in forwarding.

If one of the switches appears to flip the root from one port to another, the problem may be one of traffic prioritization (See problem five).

Another possible cause of intermittent operation is that of an auto-negotiation mismatch. If one end of the link is fixed to full-duplex mode and the peer auto-negotiates, the auto-negotiating end will fall back to half-duplex operation. At lower traffic, the volumes the link may display few if any errors. As the traffic volume rises, the fixed negotiation side will begin to experience dropped packets while the auto-negotiating side will experience collisions. Ultimately, as traffic loads approach 100%, the link will become entirely unusable. At this point, RSTP will not be able to transmit configuration messages over the link and the spanning tree topology will break down. If an alternate trunk exists, RSTP will activate it in the place of the congested port. Since activation of the alternate port often relieves the congested port of its traffic, the congested port will once again become reliable. RSTP will promptly enter it back into service, beginning the cycle once again. The root port will flip back and forth between two ports on the switch.

Problem Two

My PC/IED/Device is connected to your switch. After I reset the switch, it takes a long time before it comes up.

Is it possible that the RSTP edge setting for this port is set to false? If Edge is set to false, the bridge will make the port go through two forward delay times before the port can send or receive frames. If Edge is set to true the bridge will transition the port directly to forwarding upon link up.

Another possible explanation is that some links in the network run in half-duplex mode. RSTP uses a peer-to-peer protocol called Proposal-Agreement to ensure transitioning in the event of a link

failure. This protocol requires full-duplex operation. When RSTP detects a non-full duplex port, it cannot rely on Proposal-Agreement protocol and must make the port transition the slow (i.e. STP) way. If possible, configure the port for full-duplex operation. Otherwise, configure the port's point-to-point setting to true. Either one will allow the Proposal-Agreement protocol to be used.

Problem Three

When I test your switch by deliberately breaking a link, it takes a long time before I can poll devices past the switch. I thought RSTP was supposed to be fast. What is happening?

Is it possible that some ports participating in the topology have been configured to STP mode or that the port's point-to-point parameter is set to false? STP and multipoint ports converge slowly after failures occur.

Is it possible that the port has migrated to STP? If the port is connected to the LAN segment by shared media and STP bridges are connected to that media, then convergence after link failure will be slow.

Delays on the order of tens or hundreds of milliseconds can result in circumstances where the link broken is the sole link to the root bridge and the secondary root bridge is poorly chosen. The worst of all possible designs occurs when the secondary root bridge is located at the farthest edge of the network from the root. In this case, a configuration message will have to propagate out to the edge and then back in order to reestablish the topology.

Problem Four

My network is composed of a ring of bridges, of which two (connected to each other) are managed and the rest are unmanaged. Why does the RSTP protocol work quickly when I break a link between the managed bridges but not in the unmanaged bridge part of the ring?

A properly operating unmanaged bridge is transparent to STP configuration messages. The managed bridges will exchange configuration messages through the unmanaged bridge part of the ring as if it is non-existent. When a link in the unmanaged part of the ring fails however, the managed bridges will only be able to detect the failure through timing out of hello messages. Full connectivity will require three hello times plus two forwarding times to be restored.

Problem Five

The switch is up and running and working fine. Then I start a certain application and the network becomes unstable. After I stop the application, the network goes back to running normally.

RSTP sends its configuration messages using the highest possible priority level. If CoS is configured to allow traffic flows at the highest priority level and these traffic flows burst continuously to 100% of the line bandwidth, STP may be disrupted. It is therefore advised not to use the highest CoS.

Problem Six

After I bring up a new port, the root moves on to that port, and I don't want it to. The port that I want to become root won't do so.

Is it possible that the port cost is incorrectly programmed or that auto-negotiation derives an undesired value? Inspect the port and path costs with each port active as root.

Problem Seven

My IED/Controller doesn't work with your switch.

Certain low CPU bandwidth controllers have been found to behave less than perfectly when they receive unexpected traffic. Try disabling STP for the port.

If the controller fails around the time of a link outage then there is the remote possibility that frame disordering or duplication may be the cause of the problem. Try setting the root port of the failing controller's bridge to STP.

Problem Eight

My network runs fine with your switch but I occasionally lose polls to my devices.

Inspect network statistics to determine whether the root bridge is receiving TCNs around the time of observed frame loss. It may be possible that you have problems with intermittent links in your network.

Problem Nine

I'm getting a lot of TCNs at the root, where are they coming from?

Examine the RSTP port statistics to determine the port from which the TCNs are arriving. Sign-on to the switch at the other end of the link attached to that port. Repeat this step until the switch generating the TCNs is found (i.e. the switch that is itself not receiving a large number of TCNs). Determine the problem at that switch.

6. VLANs

ROS® provides the following VLAN features:

- Support for up to 255 VLANs
- Configurable port-native VLAN.
- Port modes of operation tailored to edge devices (such as a PC or IED) and to network switch interconnections.
- A default setting that ensures configuration-free connectivity in certain scenarios.
- Ability to force either tagged or untagged operation on the port-native VLAN.
- Ability to switch between VLAN-aware and VLAN-unaware modes of operation.
- GARP VLAN Registration Protocol (GVRP).
- Double VLAN-tagging, or QinQ
- Configurable management VLAN

6.1. VLAN Operation

6.1.1. VLANs and Tags

A virtual LAN or VLAN is a group of devices on one or more LAN segments that communicate as if they were attached to the same physical LAN segment. VLANs are extremely flexible because they are based on logical instead of physical connections.

When VLANs are introduced, all traffic in the network must belong to one or another VLAN. Traffic on one VLAN cannot pass to another, except through an internetwork router or Layer 3 switch.

A VLAN tag is the identification information that is present in frames in order to support VLAN operation.

6.1.2. Tagged vs. Untagged Frames

Tagged frames are frames with 802.1Q (VLAN) tags that specify a valid VLAN identifier (VID). Untagged frames are frames without tags or frames that carry 802.1p (Prioritization) tags only having prioritization information and a VID of 0. Frames with a VID=0 are also called priority-tagged frames.

When a switch receives a tagged frame, it extracts the VID and forwards the frame to other ports in the same VLAN.

6.1.3. Native VLAN

Each port is assigned a native VLAN number, the Port VLAN ID (PVID). When an untagged frame ingresses a port, it is associated with the port's native VLAN.

By default, when the switch transmits a frame on the native VLAN, it sends the frame untagged. The switch can be configured to transmit frames on the native VLAN tagged.

6.1.4. Management VLAN

Management traffic, like all traffic on the network, must belong to a specific VLAN. The management VLAN is configurable and always defaults to VLAN 1. This VLAN is also the default native VLAN for all ports, thus allowing all ports the possibility of managing the product. Changing the management VLAN can be used to restrict management access to a specific set of users.

6.1.5. Edge and Trunk Port Types

Each port can be configured to take on a type of Edge or Trunk.

Edge Type

An Edge port attaches to a single end device (such as a PC or IED) and carries traffic on a single pre-configured VLAN, the native VLAN.

Trunk Type

Trunk ports are part of the network and carry traffic for all VLANs between switches.

Trunk ports are automatically members of all VLANs configured in the switch.

The switch can “pass through” traffic, forwarding frames received on one trunk port out another trunk port. The trunk ports must be members of all the VLANs the “pass through” traffic is part of, even if none of those VLANs are used on edge ports.

Frames transmitted out of the port on all VLANs other than the port’s native VLAN are always sent tagged.

Note

Sometimes it may be desirable to manually restrict the traffic on the trunk to a certain group of VLANs. For example, when the trunk connects to a device (such as a Layer 3 router) that supports a subset of the available VLANs. The trunk port can be prevented from being a member of the VLAN by including it in the VLAN’s Forbidden Ports list.

Port Type	VLANs Supported	PVID Format	Usage
Edge	1 (Native) Configured	Untagged	**VLAN Unaware networks** – All frames are sent and received without the need for VLAN tags.
		Tagged	**VLAN Aware networks** – VLAN traffic domains are enforced on a single VLAN.
Trunk	All Configured	Tagged or Untagged	**Switch-to-Switch connections** – VLANs must be manually created and administered or can be dynamically learned through GVRP. **Multiple-VLAN end devices** – Implement connections to end devices that support multiple VLANs at the same time.

6.1.6. VLAN Ingress and Egress Rules

Ingress Rules

These are the VLAN ingress rules, i.e. the rules applied to all frames when they are received by the switch:

Frame received This does not depend on ingress port's VLAN configuration parameters	Untagged	Priority Tagged (VID=0)	Tagged (valid VID)
VLAN ID associated with the frame	PVID	PVID	VID in the tag
Frame dropped due to its tagged/untagged format	No	No	No
Frame dropped, if frame associated with VLAN not configured (or learned) in the switch	N/A	N/A	Yes
Frame dropped, if ingress port is not a member of the VLAN the frame is associated with	N/A	N/A	No

Egress Rules

These are the VLAN egress rules, i.e. the rules applied to all frames when they are transmitted by the switch:

Frame sent Egress port type	On egress port's native VLAN	On other VLAN	
		Port is a member of the VLAN	Port is not a member of the VLAN
Edge	According to the egress port's "PVID Format" parameter	N/A (frame is dropped)	
Trunk		Tagged	dropped

6.1.7. Forbidden Ports List

Each VLAN can be configured to exclude ports from membership in the VLAN.

6.1.8. VLAN-aware and VLAN-unaware Operation Modes

The native operation mode for an IEEE 802.1Q compliant switch is VLAN-aware. Even if a specific network architecture does not use VLANs, ROS default VLAN settings allow the switch still to operate in a VLAN-aware mode while providing functionality required for almost any network application. However, the IEEE 802.1Q standard defines a set of rules that must be followed by all VLAN-aware switches, for example:

- Valid VID range is 1 to 4094 (VID=0 and VID=4095 are invalid).
- Each frame ingressing a VLAN-aware switch is associated with a valid VID.
- Each frame egressing a VLAN-aware switch is either untagged or tagged with a valid VID (this means priority-tagged frames with VID=0 are never sent out by a VLAN-aware switch).

It turns out that some applications have requirements conflicting with the IEEE 802.1Q native mode of operation (e.g. some applications explicitly require priority-tagged frames to be received by end devices).

To ensure the required operation in any possible application scenario and provide full compatibility with legacy (VLAN-unaware) devices, RuggedSwitch® can be configured to work in a VLAN-unaware mode.

In that mode:

- Frames ingressing a VLAN-unaware switch are not associated with any VLAN.
- Frames egressing a VLAN-unaware switch are sent out unmodified, i.e. in the same untagged, 802.1Q-tagged or priority-tagged format as they were received.

6.1.9. GVRP (GARP VLAN Registration Protocol)

GVRP is a standard protocol built on GARP (the Generic Attribute Registration Protocol) to automatically distribute VLAN configuration information in a network. Each switch in a network needs only to be configured with VLANs it requires locally; it dynamically learns the rest of the VLANs configured elsewhere in the network via GVRP. A GVRP-aware end station, configured for a particular VLAN ID, can be connected to a trunk on a GVRP-aware switch and automatically become part of the desired VLAN.

When a switch sends GVRP BPDUs out of all GVRP-enabled ports, GVRP BPDUs advertise all the VLANs known to that switch (configured manually or learned dynamically through GVRP) to the rest of the network.

When a GVRP-enabled switch receives a GVRP BPDU advertising a set of VLANs, the receiving port becomes a member of those advertised VLANs and the switch begins advertising those VLANs via all the GVRP-enabled ports (other than the port on which the VLANs were learned).

To improve network security using VLANs, GVRP-enabled ports may be configured to prohibit the learning of any new dynamic VLANs but at the same time be allowed to advertise the VLANs configured on the switch.

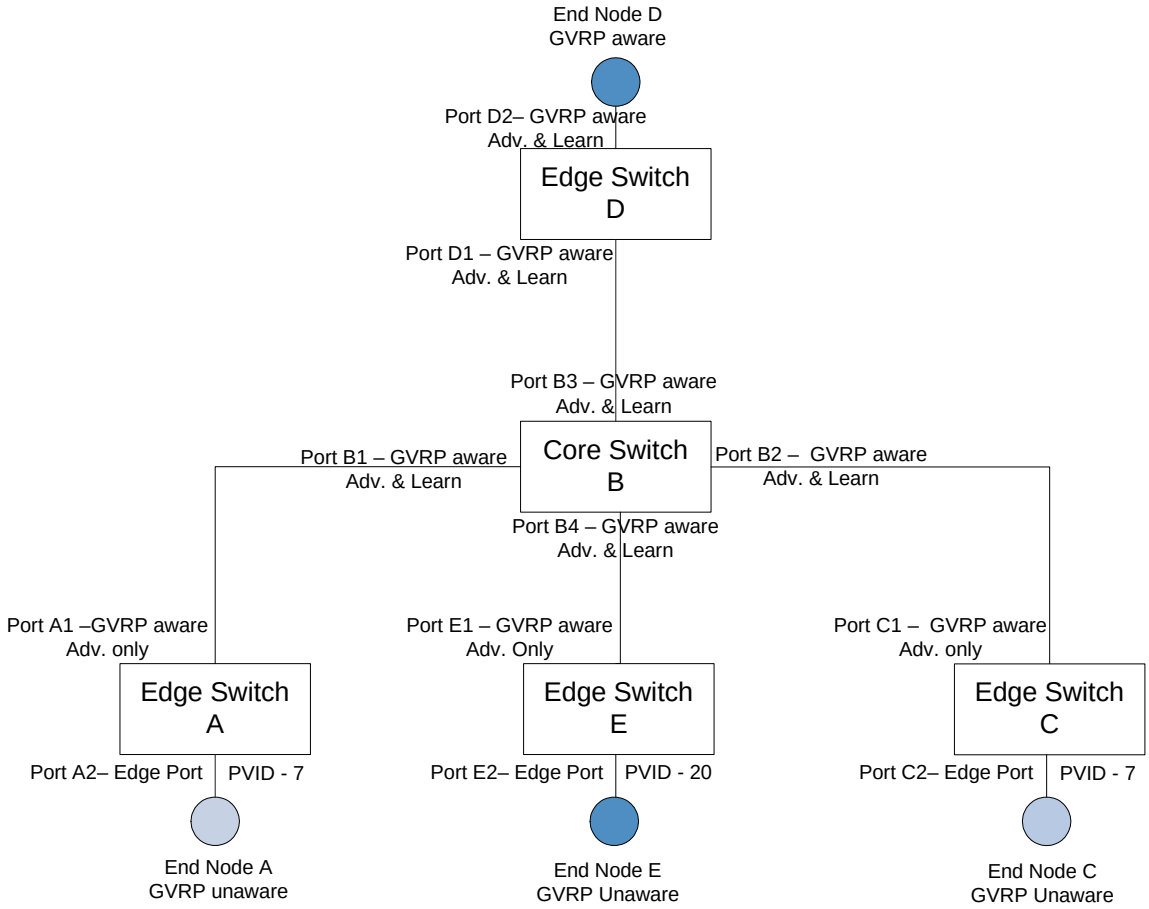

Figure 6.1. Using GVRP

An example of using GVRP:

- Ports A2, and C2 are configured with PVID 7 and port E2 is configured with PVID 20.
- End Node D is GVRP aware and is interested in VLAN 20, hence VLAN 20 is advertised by it towards switch D.
- D2 becomes member of VLAN 20.
- Ports A1 and C1 advertise VID 7 and ports B1 and B2 become members of VLAN 7.
- Ports D1 and B1 advertise VID 20 and ports B3, B4 and D1 become members of VLAN 20.

6.1.10. PVLAN Edge

PVLAN Edge (Protected VLAN Edge port) refers to a feature of the switch whereby multiple VLAN Edge ports on a single device are effectively isolated from one another. All VLAN Edge ports in a switch that are configured as "protected" in this way are prohibited from sending frames to each other, but are still allowed to send frames to other, non-protected, ports within the same VLAN. This protection extends to all traffic on the VLAN: unicast, multicast, or broadcast.

Note that this feature is strictly local to the switch. PVLAN Edge ports are not prevented from communicating with ports off the switch, whether protected (remotely) or not.

6.1.11. QinQ

QinQ is also known as double VLAN-tagging or as Nested VLANs. It is used to overlay a private Layer 2 network over a public Layer 2 network.

A large network service provider, for example, might have several clients whose networks each use multiple VLANs. It is likely that the VLAN IDs used by these different client networks would conflict with one another, were they mixed together in the provider's network. Using double VLAN-tagging, each client network could be further tagged using a client-specific VID at the edges where the clients' networks are connected to the network service provider's infrastructure.

Frames ingressing an edge port of the service provider switch are tagged with VIDs of the customer's private network. When those frames egress the switch's QinQ-enabled port into the service provider network the switch always adds an extra tag (called outer tag) on top of the frames' original VLAN tag (called inner tag) and the outer tag VID is the PVID of the frames' ingress edge port. This means that traffic from an individual customer is tagged with his unique VID and is thus segregated from other customers' traffic.

Within the service provider network, switching is based on the VID in the outer tag.

When double-tagged frames leave the service provider network, they egress a QinQ-enabled port of another switch. The switch strips the outer tag while associating the frames with the VID extracted from it before stripping. Thus, the frames are switched to appropriate edge ports, i.e. to appropriate customers.

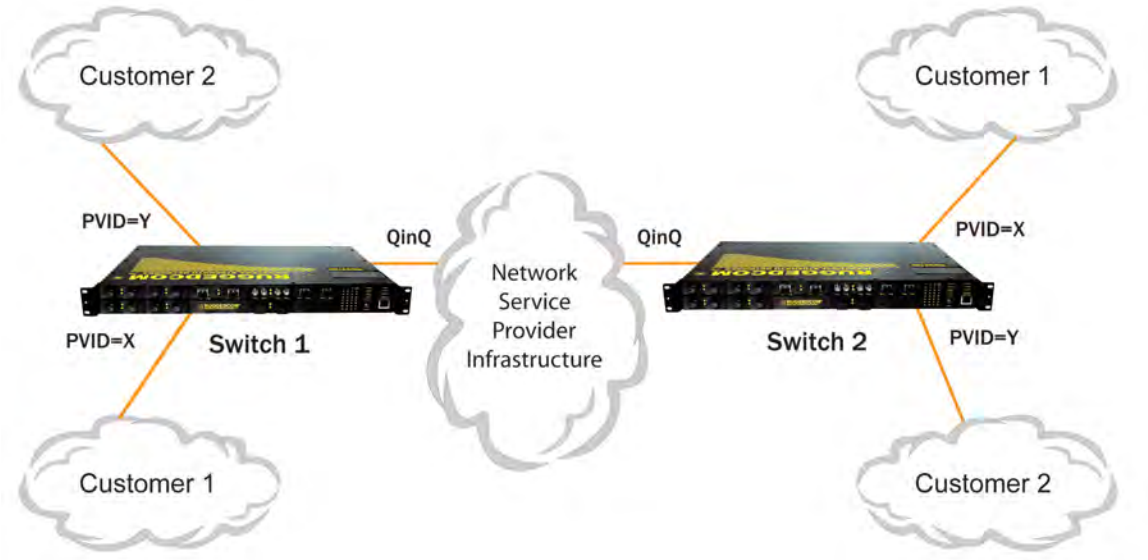

Figure 6.2. Using QinQ Example

Note

QinQ can only be enabled on one switch port at a time.

Note

Some RuggedSwitch models only support QinQ if all edge ports are configured with the same PVID. In this case, a dedicated switch must be assigned to each customer.

6.2. VLAN Applications

6.2.1. Traffic Domain Isolation

VLANs are most often used for their ability to restrict traffic flows between groups of devices.

Unnecessary broadcast traffic can be restricted to the VLAN that requires it. Broadcast storms in one VLAN need not affect users in other VLANs.

Hosts on one VLAN can be prevented from accidentally or deliberately assuming the IP address of a host on another VLAN.

By configuring the management VLAN, a management domain can be established that restricts the number of users able to modify the configuration of the network.

The use of creative bridge filtering and multiple VLANs can carve seemingly unified IP subnets into multiple regions policed by different security/access policies.

Multi-VLAN hosts can assign different traffic types to different VLANs.

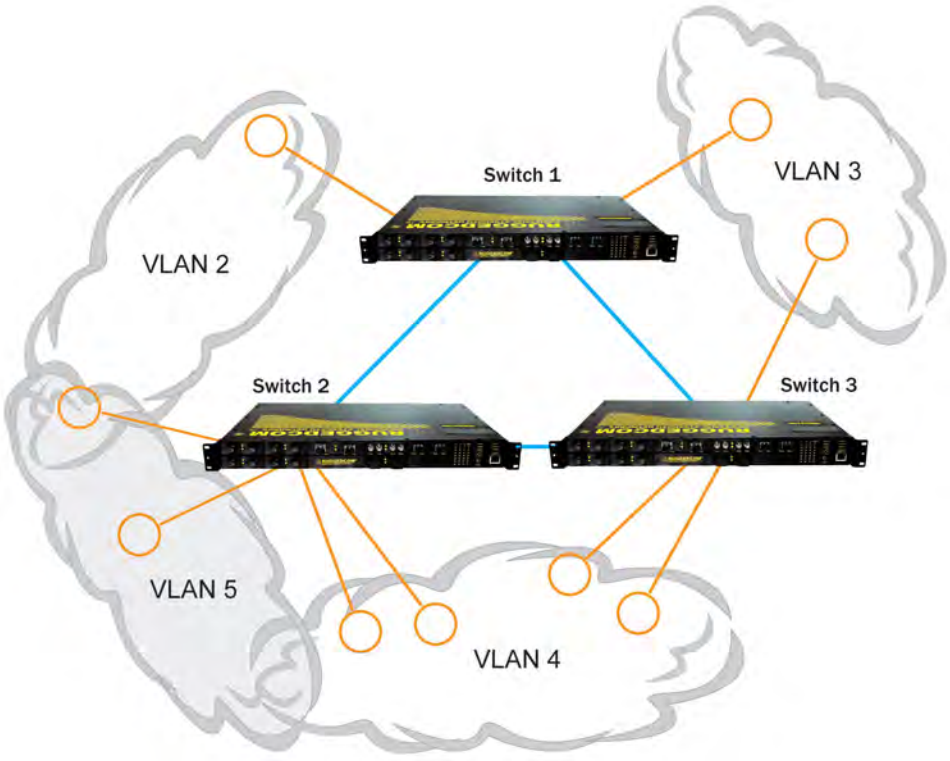

Figure 6.3. Multiple Overlapping VLANs

6.2.2. Administrative Convenience

VLANs enable equipment moves to be handled by software reconfiguration instead of by physical cable management. When a host's physical location is changed, its connection point is often changed as well. With VLANs, the host's VLAN membership and priority are simply copied to the new port.

6.2.3. Reduced Hardware

Without VLANs, traffic domain isolation requires using separate bridges for separate networks. VLANs eliminate the need for separate bridges.

The number of network hosts may often be reduced. Often, a server is assigned to provide services for independent networks. These hosts may be replaced by a single, multi-homed host supporting each network on its own VLAN. This host can perform routing between VLANs.

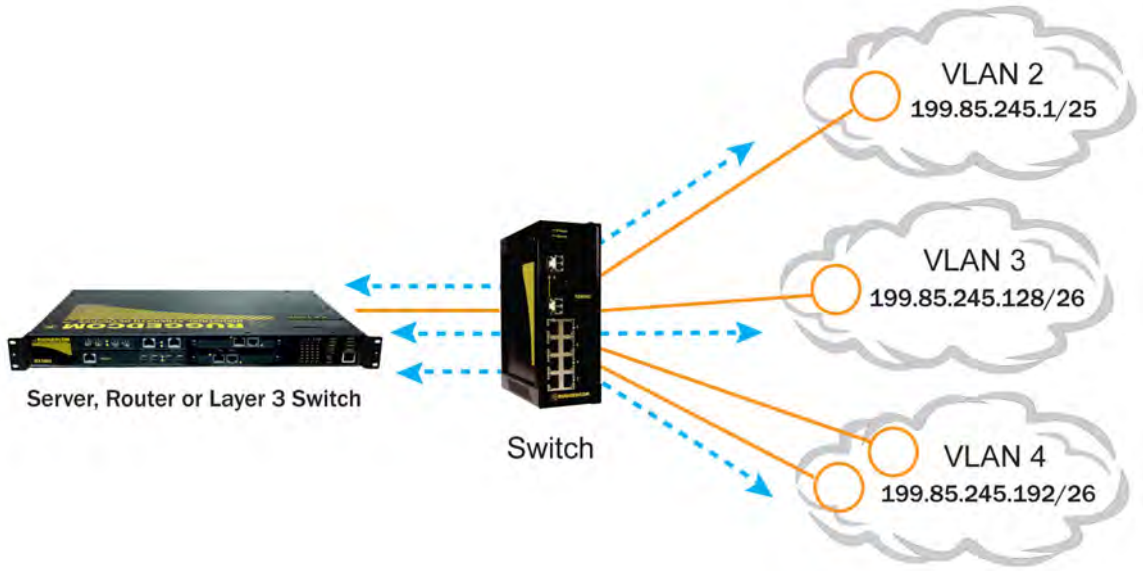

Figure 6.4. Inter-VLAN Communications

6.3. VLAN Configuration

The Virtual LANs menu is accessible from the main menu.

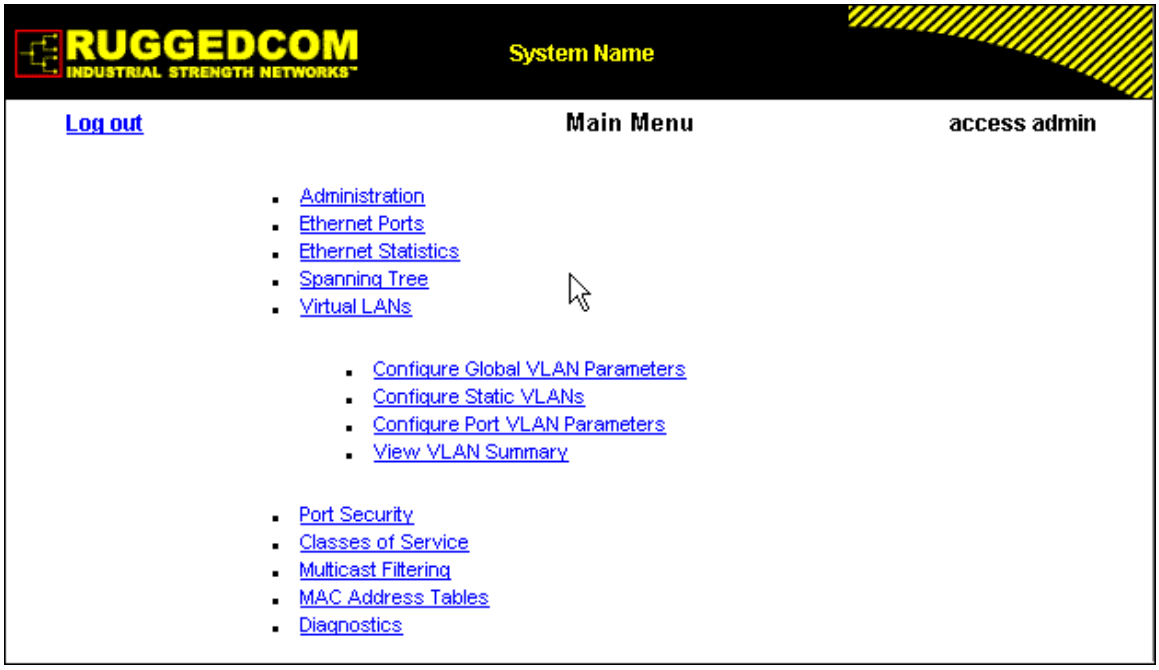

Figure 6.5. Virtual LANs Menu

6.3.1. Global VLAN Parameters

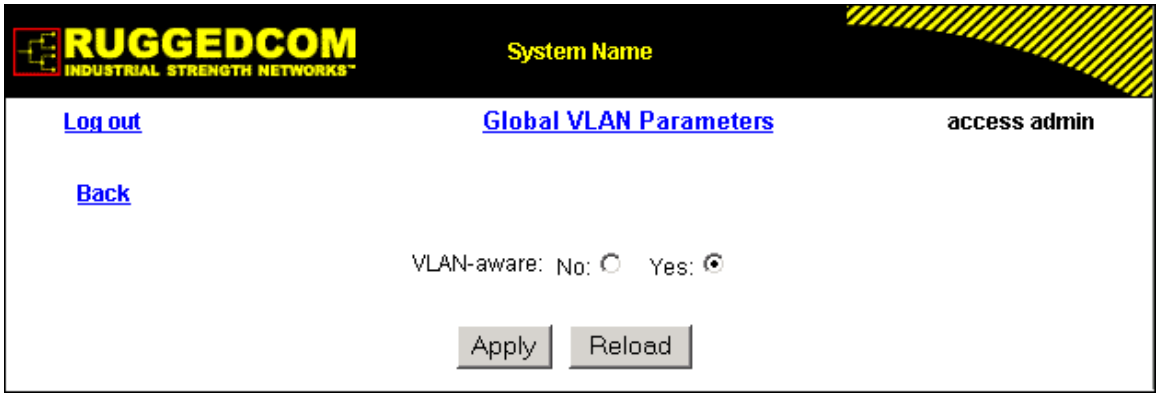

Figure 6.6. Global VLAN Parameters Form

VLAN-aware

```
Synopsis: { No, Yes }
Default: Yes
```

Set either VLAN-aware or VLAN-unaware mode of operation.

Note

Do not attempt to change the "VLAN-aware" parameter of the managed switch by applying a configuration (.CSV) file update. Configuration file updates are used to apply "bulk changes" to the current configuration of a switch. Instead, a change to this individual parameter MUST first be applied separately from any other table (i.e. parameter) changes. In other words, configuration file updates should exclude the "VLAN-aware" parameter.

6.3.2. Static VLANs

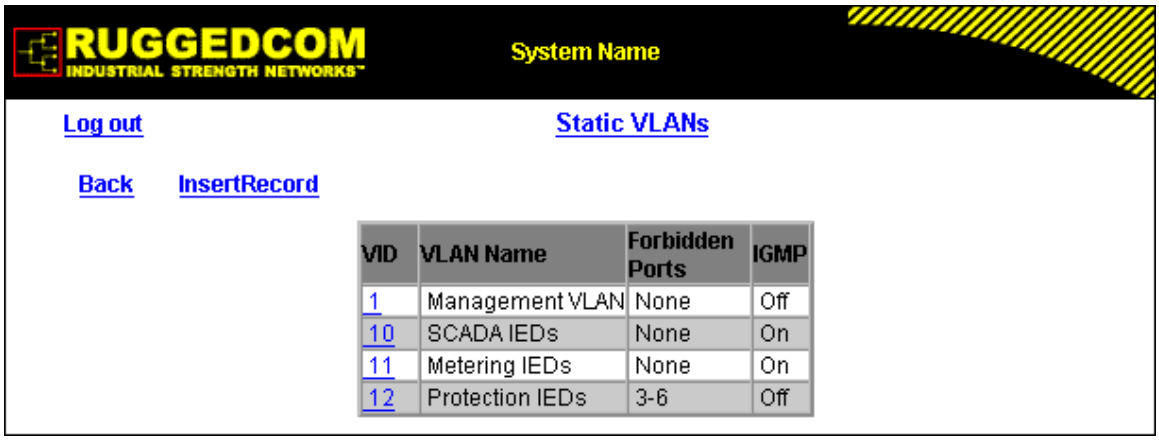

Figure 6.7. Static VLANs Table

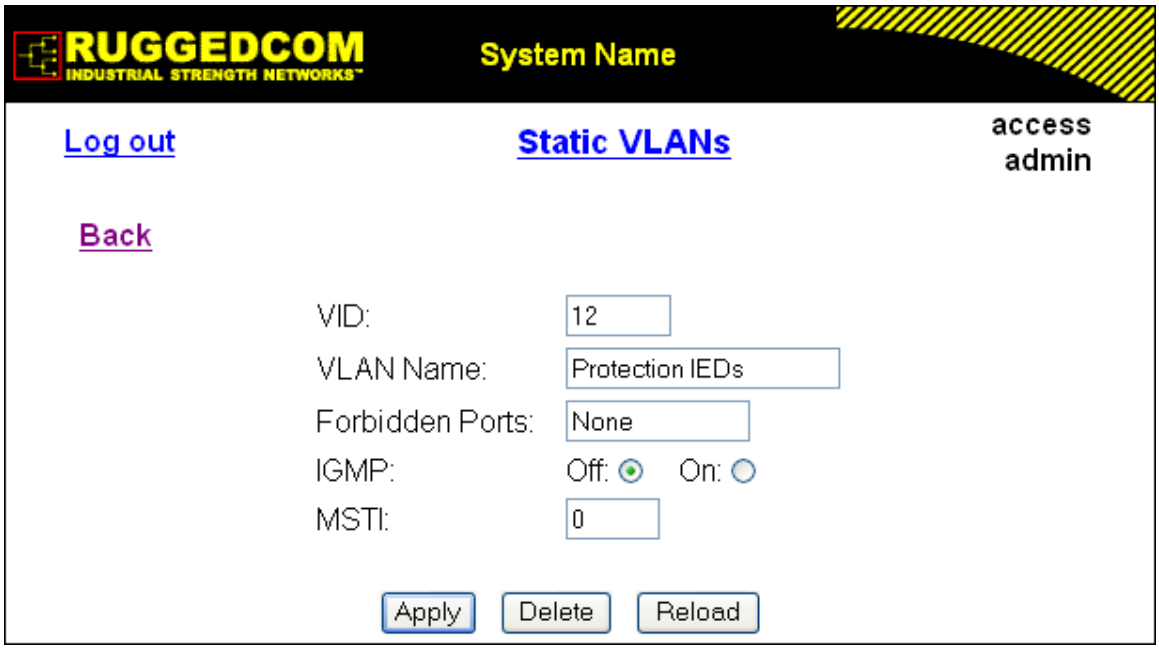

Figure 6.8. Static VLANs Form

VID

```
Synopsis: 1 to 4094
Default: 1
```

The VLAN Identifier is used to identify the VLAN in tagged Ethernet frames according to IEEE 802.1Q.

VLAN Name

```
Synopsis: Any 19 characters
Default:
```

The VLAN name provides a description of the VLAN purpose (for example, Engineering VLAN).

Forbidden Ports

```
Synopsis: Any combination of numbers valid for this parameter
Default: None
```

These are ports that are not allowed to be members of the VLAN.

Examples:

None - all ports of the switch are allowed to be members of the VLAN

2,4-6,8 - all ports except ports 2,4,5,6 and 8 are allowed to be members of the VLAN

IGMP

```
Synopsis: { Off, On }
Default: Off
```

This parameter enables or disables IGMP Snooping on the VLAN.

MSTI

```
Synopsis: 0 to 16
Default: 0
```

This parameter is only valid for Multiple Spanning Tree Protocol (MSTP) and has no effect, if MSTP is not used. The parameter specifies the Multiple Spanning Tree Instance (MSTI) to which the VLAN should be mapped.

Note

If IGMP Snooping is not enabled for the VLAN, both IGMP messages and multicast streams will be forwarded directly to all members of the VLAN. If any one member of the VLAN joins a multicast group then ***all members of the VLAN will receive the multicast traffic****.*

6.3.3. Port VLAN Parameters

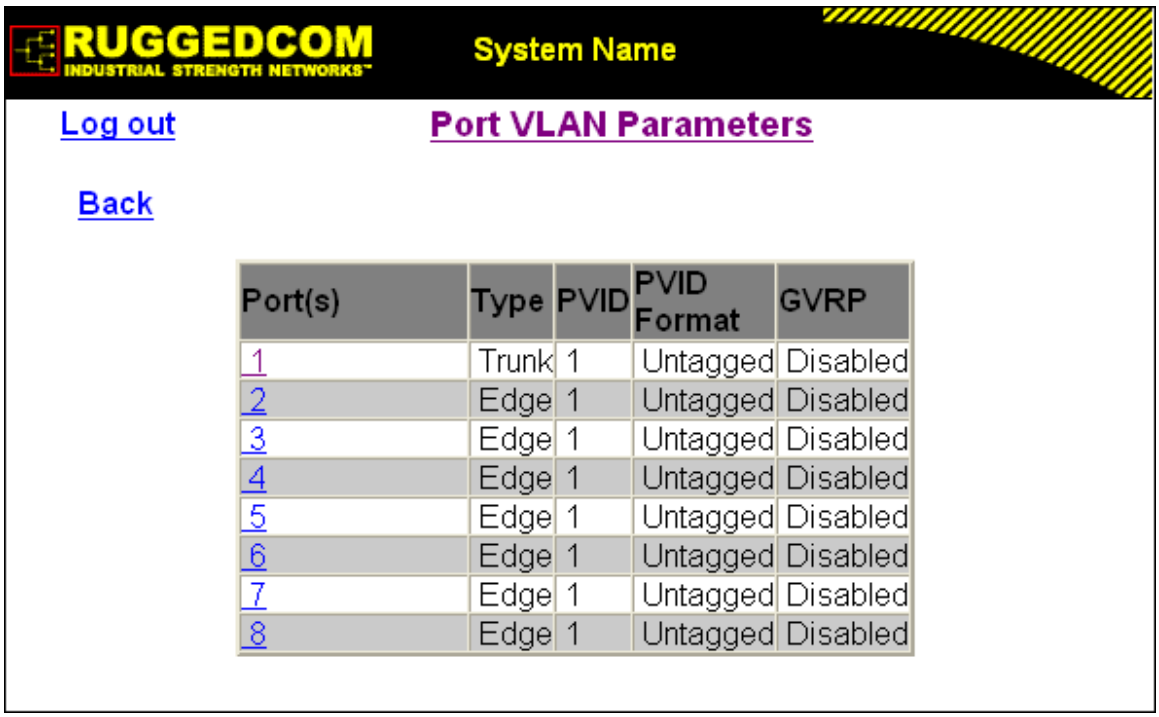

Figure 6.9. Port VLAN Parameters Table

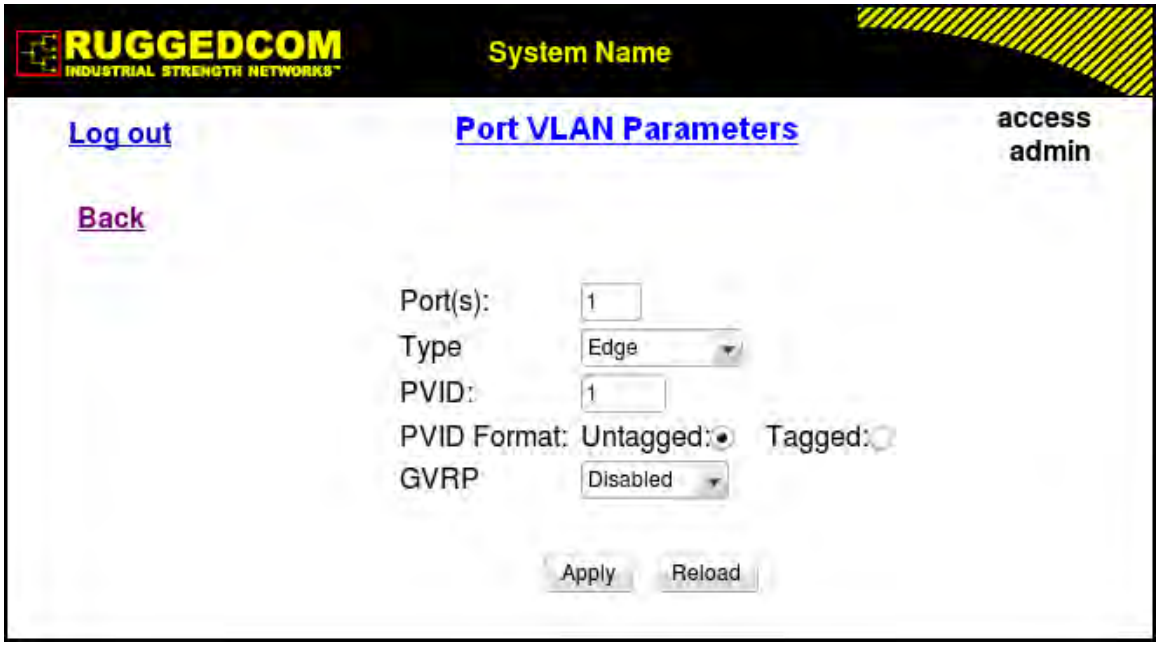

Figure 6.10. Port VLAN Parameters Form

Port(s)

`Synopsis: Any combination of numbers valid for this parameter`

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Type

`Synopsis: {Edge, Trunk, PVLANEdge, QinQ}`
`Default: Edge`

This parameter specifies how the port determines its membership in VLANs. There are few types of ports:

Edge - the port is only a member of one VLAN (its native VLAN specified by the 'PVID' parameter).

Trunk - the port is automatically a member of all configured VLANs. Frames transmitted out of the port on all VLANs except the port's native VLAN will be always tagged. It can also be configured to use GVRP for automatic VLAN configuration.

PVLANEdge - the port is only a member of one VLAN (its native VLAN specified by the 'PVID' parameter), and does not forward traffic to other PVLANedge ports within the same VLAN.

QinQ - the port is a trunk port using double-VLAN tagging, or nested VLANs. An extra VLAN tag is always added to all frames egressing this port. VID in the added extra tag is the PVID of the frame's ingress port. VLAN tag is always stripped from frames ingressing this port.

PVID

`Synopsis: 1 to 4094`
`Default: 1`

The Port VLAN Identifier specifies the VLAN ID associated with untagged (and 802.1p priority tagged) frames received on this port.

Frames tagged with a non-zero VLAN ID will always be associated with the VLAN ID retrieved from the frame tag.

Modify this parameter with care! By default, the switch is programmed to use VLAN 1 for management and every port on the switch is programmed to use VLAN 1. If you modify a switch port to use a VLAN other than the management VLAN, devices on that port will not be able to manage the switch.

PVID Format

```
Synopsis: { Untagged, Tagged }
Default: Untagged
```

Specifies whether frames transmitted out of the port on its native VLAN (specified by the 'PVID' parameter) will be tagged or untagged.

GVRP

```
Synopsis: { Adv&Learn, Adv Only, Disabled }
Default: Disabled
```

Configures GVRP (Generic VLAN Registration Protocol) operation on the port. There are several GVRP operation modes:

DISABLED - the port is not capable of any GVRP processing.

ADVERTISE ONLY - the port will declare all VLANs existing in the switch (configured or learned) but will not learn any VLANs.

ADVERTISE & LEARN - the port will declare all VLANs existing in the switch (configured or learned) and can dynamically learn VLANs.

Only Trunk ports are GVRP-capable.

6.3.4. VLAN Summary

There are actually three ways that a VLAN can be created in the switch:

Explicit

A VLAN is explicitly configured in the Static VLANs list.

Implicit

A VLAN ID is a parameter required for different feature configurations (e.g. Port VLAN Parameters, Static MAC Addresses, IP Interface Type and ID). When such a parameter is set to some VLAN ID value, appropriate VLAN is automatically created, if it does not yet exist.

Dynamic

A VLAN learned through GVRP.

Note

Not explicitly created VLAN is always created with IGMP Snooping disabled. If it is desirable for IGMP to be used on that VLAN, it should be created as a Static VLAN with IGMP enabled.

All VLANs, regardless of the way they were created, are shown in the VLAN Summary.

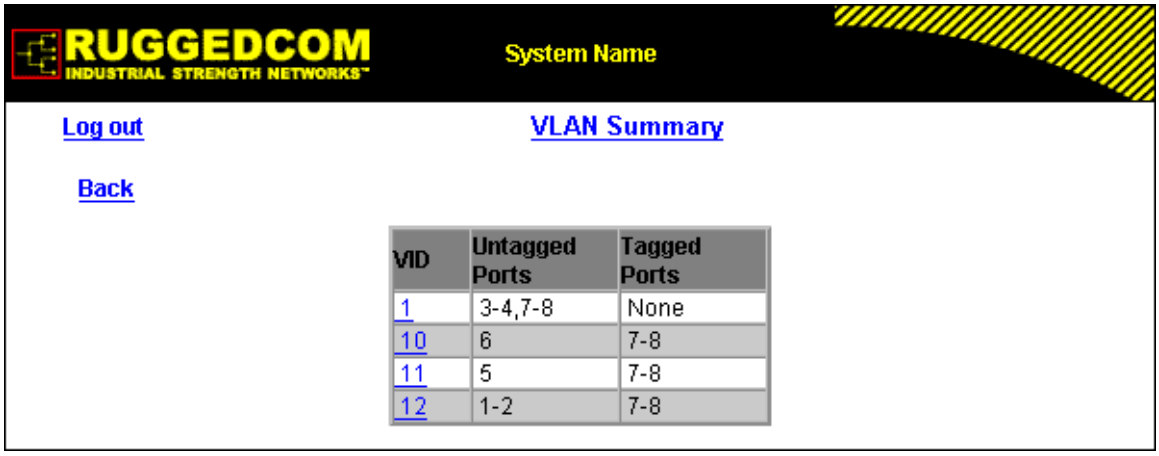

Figure 6.11. VLAN Summary Table

VID

`Synopsis: 1 to 4094`

The VLAN Identifier is used to identify the VLAN in tagged Ethernet frames according to IEEE 802.1Q.

Untagged Ports

`Synopsis: Any combination of numbers valid for this parameter`

All ports that are untagged members of the VLAN.

Tagged Ports

`Synopsis: Any combination of numbers valid for this parameter`

All ports that are tagged members of the VLAN.

6.4. Troubleshooting

Problem One

I don't need VLANs at all. How do I turn them off?

Simply leave all ports set to type "Edge" and leave the native VLAN set to 1. This is the default configuration for the switch.

Problem Two

I have added two VLANs: 2 and 3. I made a number of ports members of these VLANS. Now I need some of the devices in one VLAN to send messages to some devices in the other VLAN.

If the devices need to communicate at the physical address layer, they must be members of the same VLAN. If they can communicate in a Layer 3 fashion (i.e. using a protocol such as IP or IPX), you can use a router. The router will treat each VLAN as a separate interface, which will have its own associated IP address space.

Problem Three

I have a network of thirty switches for which I wish to restrict management traffic to a separate domain. What is the best way of doing this while still staying in contact with these switches?

At the switch where the management station is located, configure a port to use the new management VLAN as its native VLAN. Configure a host computer to act as a temporary management station.

At each switch, configure the management VLAN to the new value. As each switch is configured, you will immediately lose contact with it, but should be able to re-establish communications from the temporary management station. After all switches have been taken to the new management VLAN, configure the ports of all attached management devices to use the new VLAN.

Note

Establishing a management domain is often accompanied with the establishment of an IP subnet specifically for the managed devices.

7. Port Security

ROS® Port Security provides you with the following features:

- Authorizing network access using Static MAC Address Table.
- Authorizing network access using IEEE 802.1X authentication.
- Configuring IEEE 802.1X authentication parameters.
- Detecting port security violation attempt and performing appropriate actions.

7.1. Port Security Operation

Port Security, or Port Access Control, provides the ability to filter or accept traffic from specific MAC addresses.

Port Security works by inspecting the source MAC addresses of received frames and validating them against the list of MAC addresses authorized on the port. Unauthorized frames will be filtered and, optionally, the port that receives the frame will be shut down permanently or for a period of time. An alarm will be raised indicating the detected unauthorized MAC address.

Frames to unknown destination addresses will not be flooded through secure ports.

Note

Port security is applied at the edge of the network in order to restrict admission to specific devices. Do not apply port security on core switch connections.

ROS supports several MAC address authorization methods described below:

7.1.1. Static MAC address-based authorization

- With this method, the switch validates the source MAC addresses of received frames against the contents in the Static MAC Address Table.
- Also, ROS supports a highly flexible Port Security configuration which provides a convenient means for network administrators to use the feature in various network scenarios.
- A Static MAC address can be configured without a port number being explicitly specified. In this case, the configured MAC address will be automatically authorized on the port where it is detected. This allows devices to be connected to any secure port on the switch without requiring any reconfiguration.
- The switch can also be programmed to learn (and, thus, authorize) a pre-configured number of the first source MAC addresses encountered on a secure port. This enables the capture of the appropriate secure addresses when first configuring MAC address-based authorization on a port. Those MAC addresses are automatically inserted into the Static MAC Address Table and remain there until explicitly removed by the user.

7.1.2. IEEE 802.1X Authentication

The IEEE 802.1X standard defines a mechanism for port-based network access control and provides a means of authenticating and authorizing devices attached to LAN ports.

Although 802.1X is mostly used in wireless networks, this method is also implemented in wired switches.

The 802.1X standard defines three major components of the authentication method: Supplicant, Authenticator and Authentication server.

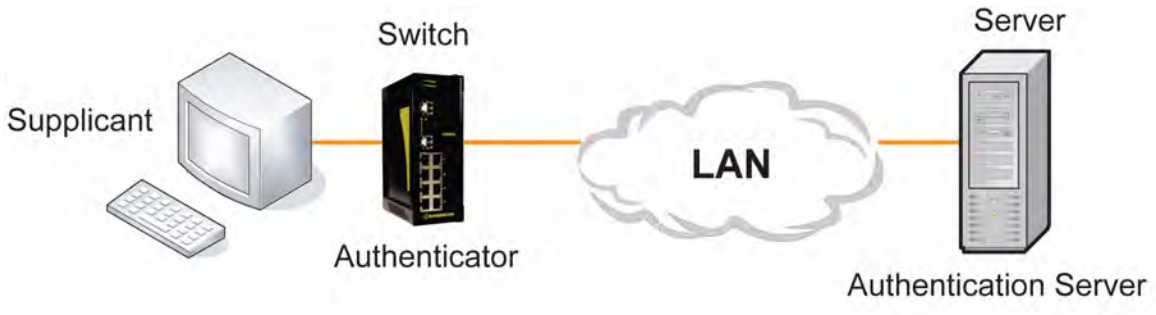

Figure 7.1. 802.1X General Topology

RuggedSwitch® supports the Authenticator component.

802.1X makes use of Extended Authentication Protocol (EAP) which is a generic PPP authentication protocol and supports various authentication methods. 802.1X defines a protocol for communication between the Supplicant and the Authenticator, EAP over LAN (EAPOL).

RuggedSwitch® communicates with the Authentication Server using EAP over RADIUS.

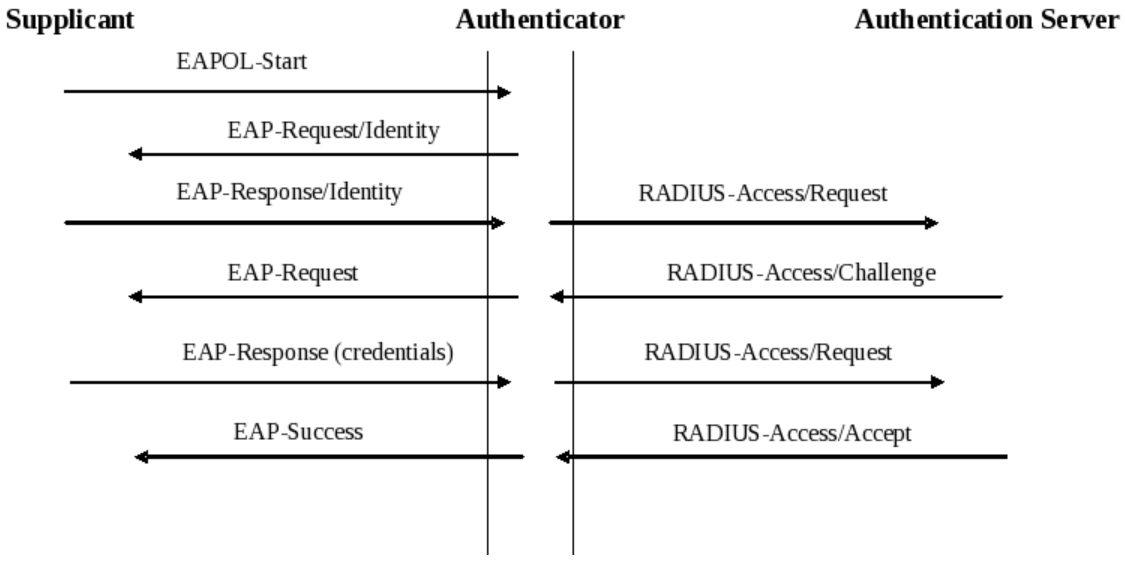

Figure 7.2. 802.1X Packet Exchange

Note

The switch supports authentication of one host per port.

Note

If the host's MAC address is configured in the Static MAC Address Table, it will be authorized, even if the host authentication is rejected by the authentication server.

7.2. Port Security Configuration

The Port Security menu is accessible from the main menu.

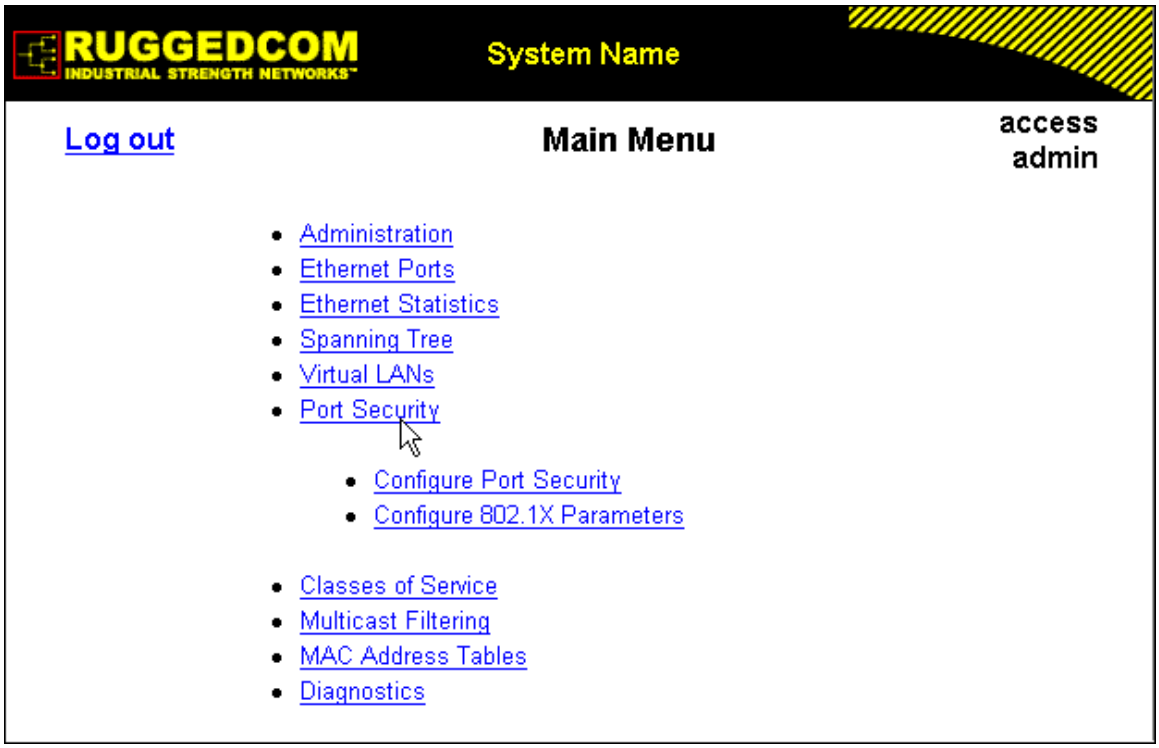

Figure 7.3. Port Security Menu

7.2.1. Port Security Parameters

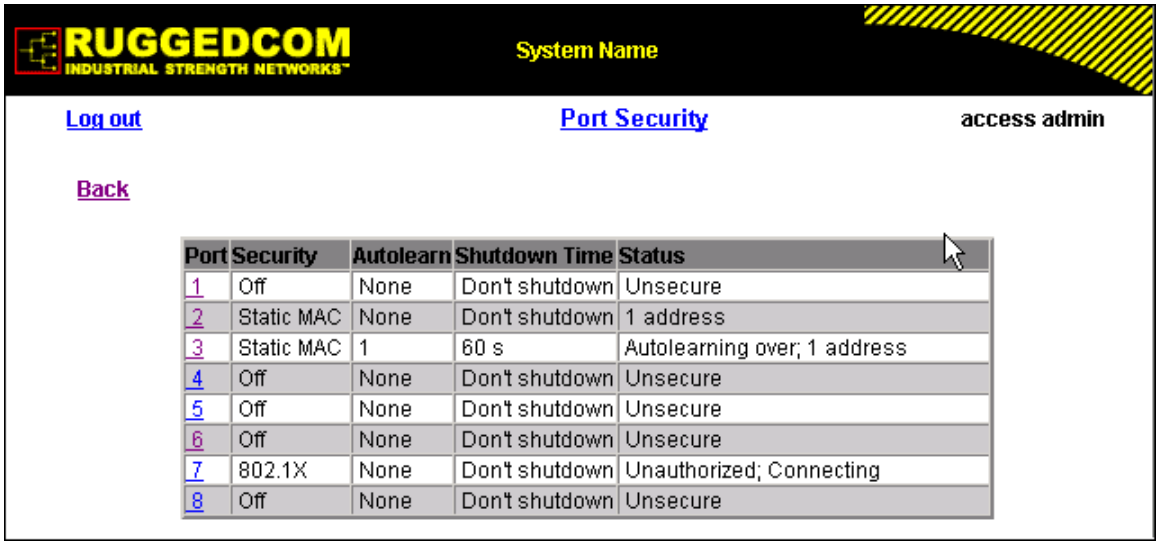

Port	Security	Autolearn	Shutdown Time	Status
1	Off	None	Don't shutdown	Unsecure
2	Static MAC	None	Don't shutdown	1 address
3	Static MAC	1	60 s	Autolearning over; 1 address
4	Off	None	Don't shutdown	Unsecure
5	Off	None	Don't shutdown	Unsecure
6	Off	None	Don't shutdown	Unsecure
7	802.1X	None	Don't shutdown	Unauthorized; Connecting
8	Off	None	Don't shutdown	Unsecure

Figure 7.4. Port Security Parameters Table

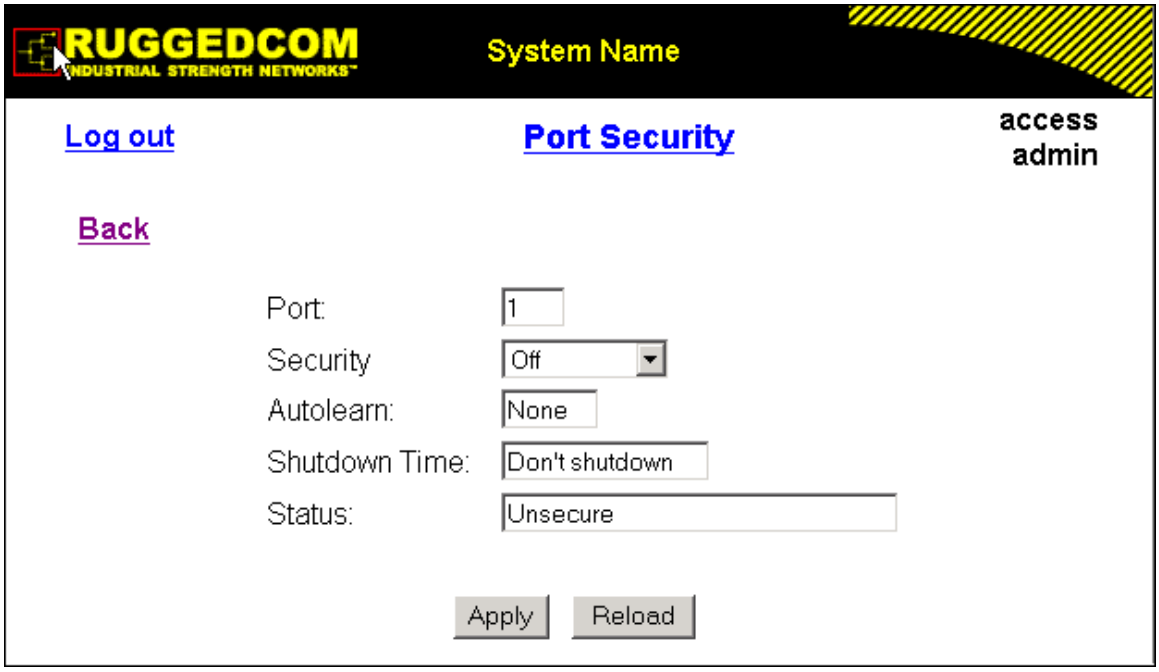

Figure 7.5. Port Security Parameters Form

Port

```
Synopsis: 1 to maximum port number
Default: 0
```

The port number as seen on the front plate silkscreen of the switch.

Security

```
Synopsis: { Off, Static MAC, 802.1X }
Default: Off
```

Enables or disables the port's security feature. Two types of port access control are available:

- Static MAC address-based. With this method, authorized MAC address(es) should be configured in the Static MAC Address table. If some MAC addresses are not known in advance (or it is not known to which port they will be connected), there is still an option to configure the switch to auto-learn certain number of MAC addresses. Once learned, they don't age out until the unit is reset or the link goes down.
- IEEE 802.1X standard authentication.

Autolearn

```
Synopsis: 1 to maximum number or { None }
Default: None
```

Specifies the maximum number of MAC addresses that can be dynamically learned on the port. If there are static addresses configured on the port, the actual number of addresses allowed to be learned is this number minus the number of the static MAC addresses.

Shutdown Time

```
Synopsis: 1 to 86400 s or { Until reset, Don't shutdown }
Default: Don't shutdown
```

Specifies for how long to shut down the port, if a security violation occurs

Status

`Synopsis: View only string of 31 characters`

Describes the security status of the port.

7.2.2. 802.1X Parameters

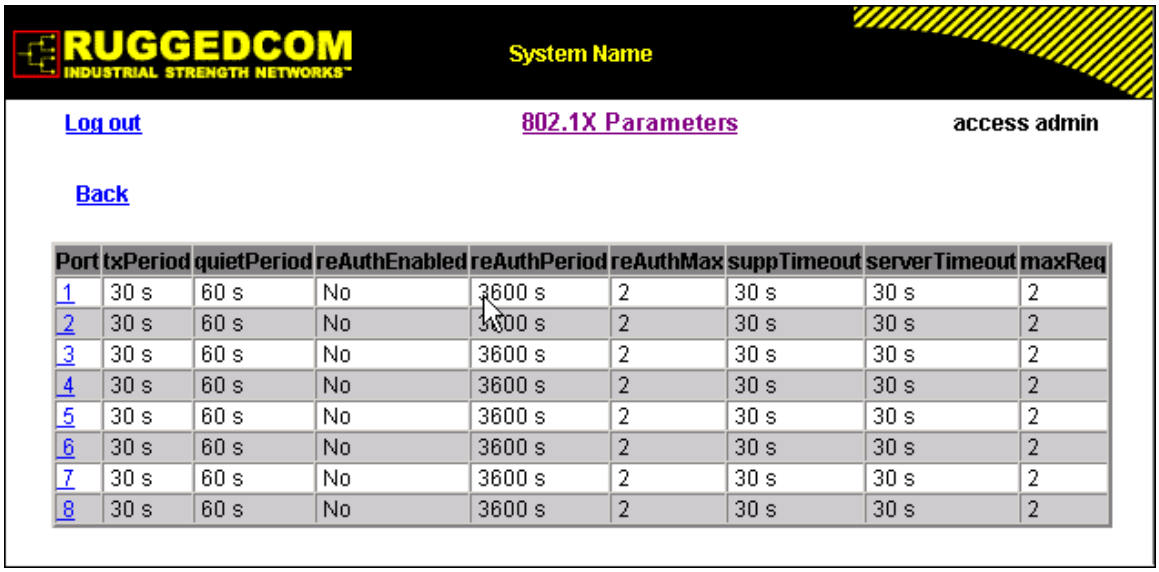

Port	txPeriod	quietPeriod	reAuthEnabled	reAuthPeriod	reAuthMax	suppTimeout	serverTimeout	maxReq
1	30 s	60 s	No	3600 s	2	30 s	30 s	2
2	30 s	60 s	No	3600 s	2	30 s	30 s	2
3	30 s	60 s	No	3600 s	2	30 s	30 s	2
4	30 s	60 s	No	3600 s	2	30 s	30 s	2
5	30 s	60 s	No	3600 s	2	30 s	30 s	2
6	30 s	60 s	No	3600 s	2	30 s	30 s	2
7	30 s	60 s	No	3600 s	2	30 s	30 s	2
8	30 s	60 s	No	3600 s	2	30 s	30 s	2

Figure 7.6. 802.1X Parameters Table

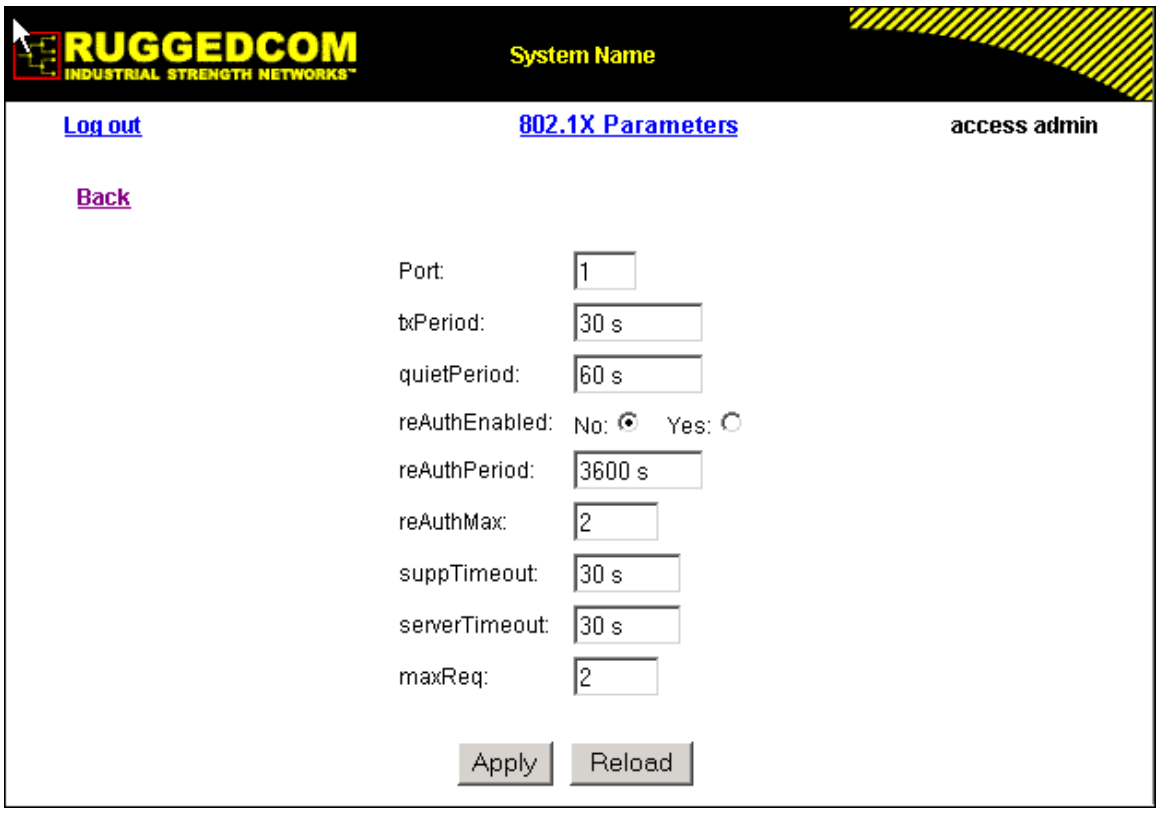

Figure 7.7. 802.1X Parameters Form

Port

```
Synopsis: 1 to maximum port number
Default: 1
```

The port number as seen on the front plate silkscreen of the switch.

txPeriod

```
Synopsis: 1 to 65535
Default: 30 s
```

The time to wait for the Supplicant's EAP Response/Identity packet before retransmitting an EAP Request/Identity packet.

quietPeriod

```
Synopsis: 0 to 65535
Default: 60 s
```

The period of time not to attempt to acquire a Supplicant after the authorization session failed.

reAuthEnabled

```
Synopsis: { No, Yes }
Default: No
```

Enables or disables periodic re-authentication.

reAuthPeriod

```
Synopsis: 60 to 86400
Default: 3600 s
```

The time between periodic re-authentication of the Supplicant.

reAuthMax

```
Synopsis: 1 to 10
Default: 2
```

The number of re-authentication attempts that are permitted before the port becomes unauthorized.

suppTimeout

```
Synopsis: 1 to 300
Default: 30 s
```

The time to wait for the Supplicant's response to the authentication server's EAP packet.

serverTimeout

```
Synopsis: 1 to 300
Default: 30 s
```

The time to wait for the authentication server's response to the Supplicant's EAP packet.

maxReq

```
Synopsis: 1 to 10
Default: 2
```

The maximum number of times to retransmit the authentication server's EAP Request packet to the Supplicant before the authentication session times out.

8. Classes of Service

ROS® CoS provides the following features:

- Support for 4 Classes of Service
- Ability to prioritize traffic by ingress port.
- Ability to prioritize traffic by the priority field in 802.1Q tags.
- Ability to prioritize traffic based on its source or destination MAC address.
- Ability to prioritize traffic by the TOS field in the IP header.

8.1. CoS Operation

CoS provides the ability to expedite the transmission of certain frames and port traffic over others.

The CoS of a frame can take on one of four values: Normal, Medium, High or Critical.

The default policies of the switch enforce a Normal CoS for all traffic.

Note

Use the highest supported CoS with caution, as it is always used by the switch for handling network management traffic such as RSTP BPDUs.

If this CoS is used for regular network traffic, upon traffic bursts, it may result in loss of some network management frames which in its turn may result in loss of connectivity over the network.

The CoS feature has two main phases - inspection and forwarding:

8.1.1. Inspection Phase

In the inspection phase, the CoS priority of a received frame is determined from:

- The priority field in 802.1Q tags.
- The Differentiated Services Code Point (DSCP) component of the Type Of Service (TOS) field, if the frame is IP.
- The default CoS for the port.
- A specific CoS based upon the source and destination MAC address (as set in the Static MAC Address Table).

Note that a frame's CoS will be determined once the first examined parameter is found in the frame.

Received frames are first examined to determine if their destination or source MAC address is found in the Static MAC Address Table. If yes, the CoS configured for the static MAC address is used. If neither destination or source MAC address is in the Static MAC Address Table, the frame

is then examined for 802.1Q tags and the priority field is mapped to a CoS. If a tag is not present, the frame is examined to determine if it is an IP frame. If the frame is IP and inspecting TOS is enabled, the CoS is determined from the DSCP field. If the frame is not IP or inspecting TOS is disabled, the default CoS for the port is used.

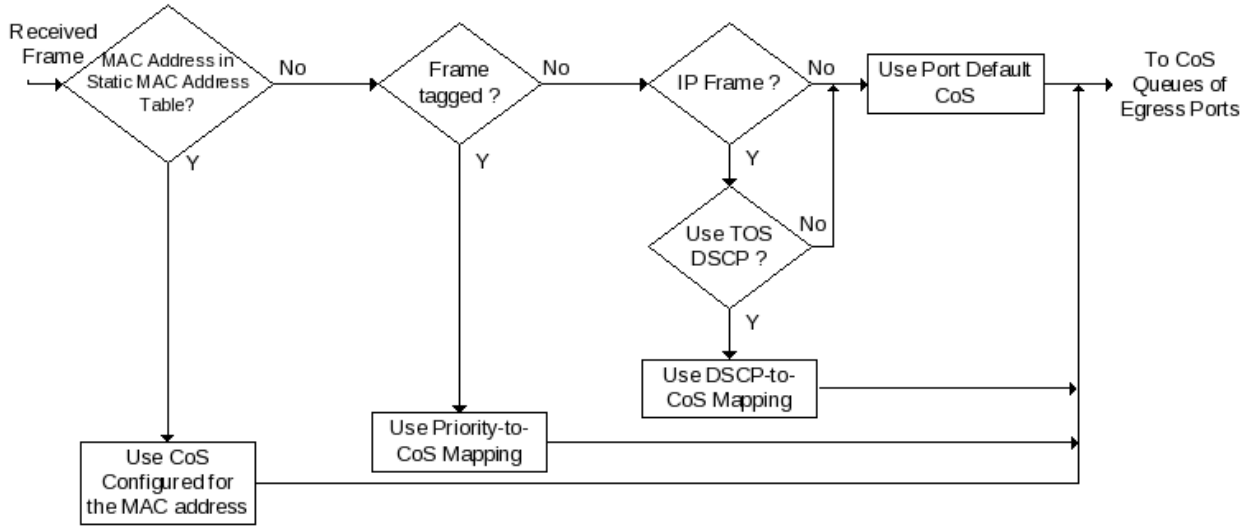

Figure 8.1. Determining The CoS Of A Received Frame

After inspection, the frame is the forwarded to the egress port for transmission.

8.1.2. Forwarding Phase

The inspection phase results in the CoS of individual frames being determined. When these frames are forwarded to the egress port, they are collected into one of the priority queues according to the CoS assigned to each frame.

CoS weighting selects the degree of preferential treatment that is attached to different priority queues. The ratio of the number of higher CoS to lower CoS frames transmitted can be programmed. If desired, the user can program lower CoS frames are to be transmitted only after all higher CoS frames have been serviced.

8.2. CoS Configuration

The Classes Of Service menu is accessible from the main menu.

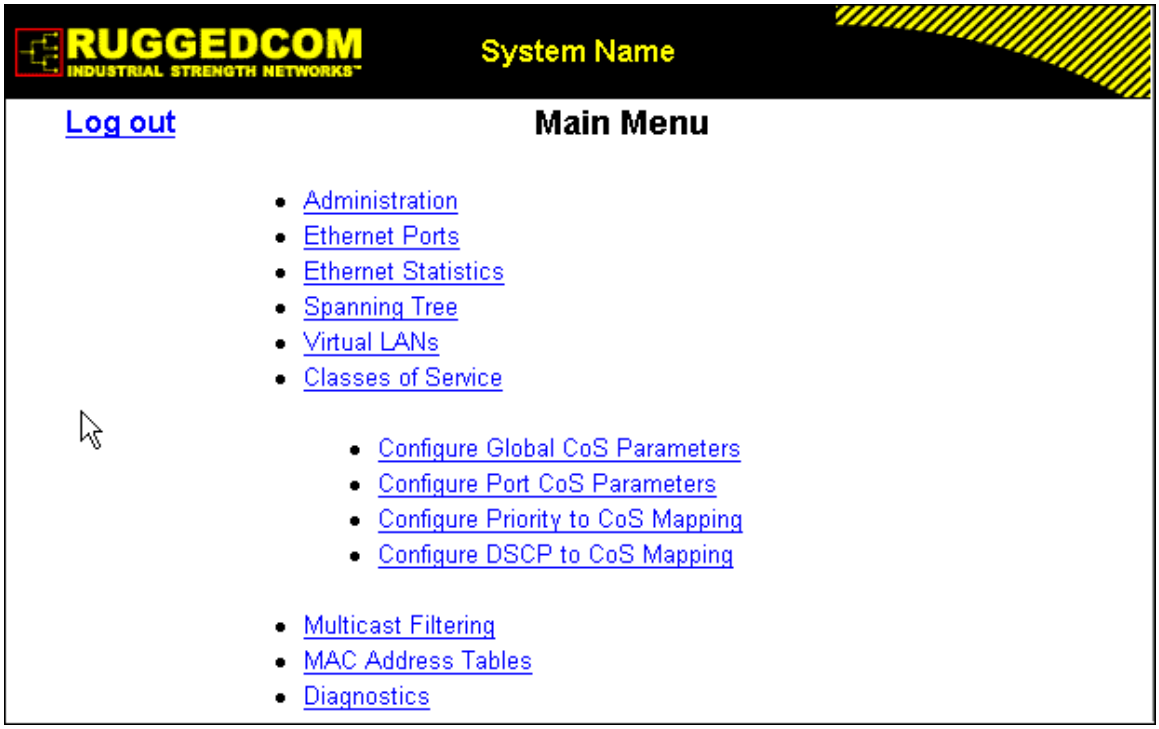

Figure 8.2. Classes Of Service Menu

8.2.1. Global CoS Parameters

Figure 8.3. Global CoS Parameters Form

CoS Weighting

```
Synopsis: { 8:4:2:1, Strict }
Default: 8:4:2:1
```

During traffic bursts, frames queued in the switch pending transmission on a port may have different CoS priorities.

This parameter specifies weighting algorithm for transmitting different priority CoS frames.

Examples:

8:4:2:1 - 8 Critical, 4 High, 2 Medium and 1 Normal priority CoS frame

Strict - lower priority CoS frames will be only transmitted after all higher priority CoS frames have been transmitted.

8.2.2. Port CoS Parameters

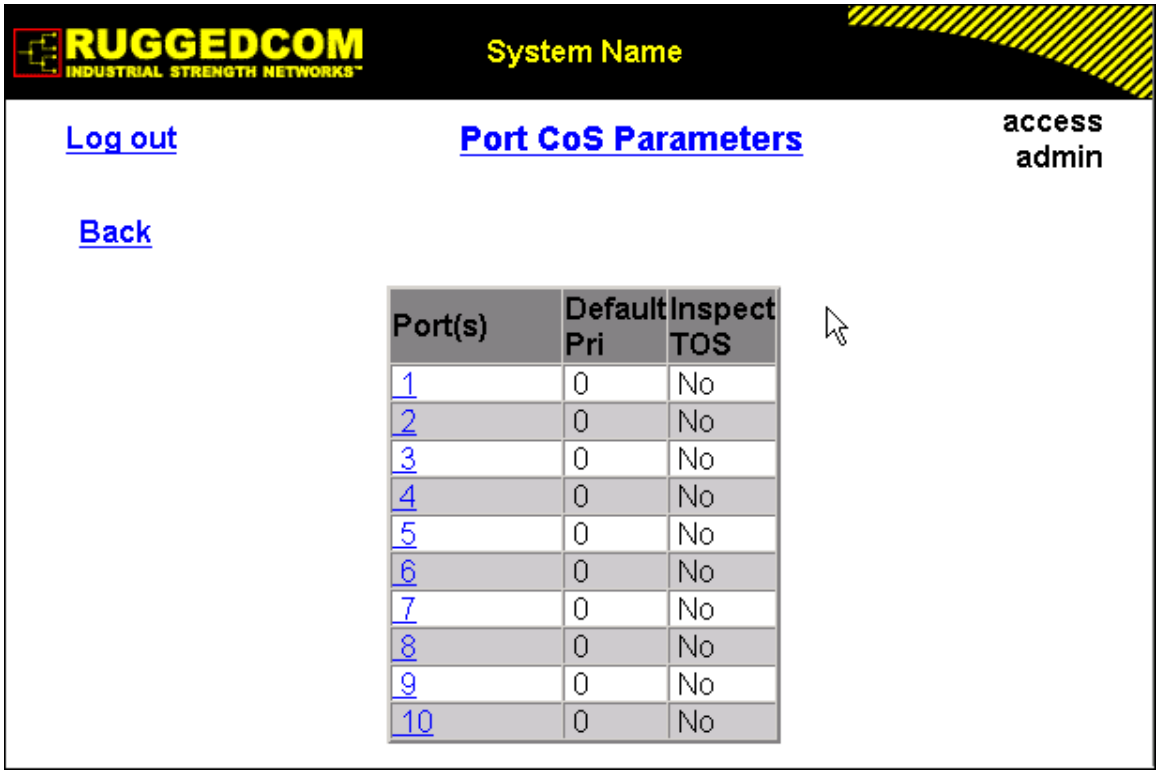

Figure 8.4. Port CoS Parameter Form

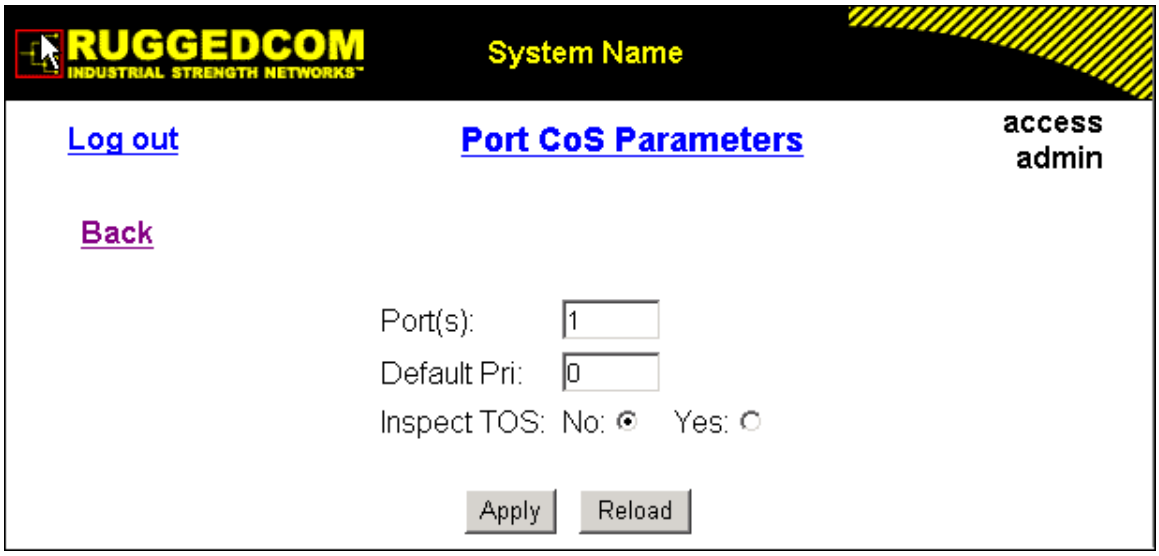

Figure 8.5. Port CoS Parameter Form

Port(s)

```
Synopsis: 1 to maximum port number
```

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Default Pri

```
Synopsis: 0 to 7
Default: 0
```

This parameter allows prioritization of the frames received on this port that are not prioritized based on the frames' contents (e.g. priority field in the VLAN tag, DiffServ field in the IP header, prioritized MAC address).

Inspect TOS

```
Synopsis: { No, Yes }
Default: No
```

This parameter enables or disables parsing of the Type-Of-Service (TOS) field in the IP header of the received frames to determine the Class of Service that should be assigned. When TOS parsing is enabled, the switch will use the Differentiated Services bits in the TOS field.

8.2.3. Priority to CoS Mapping

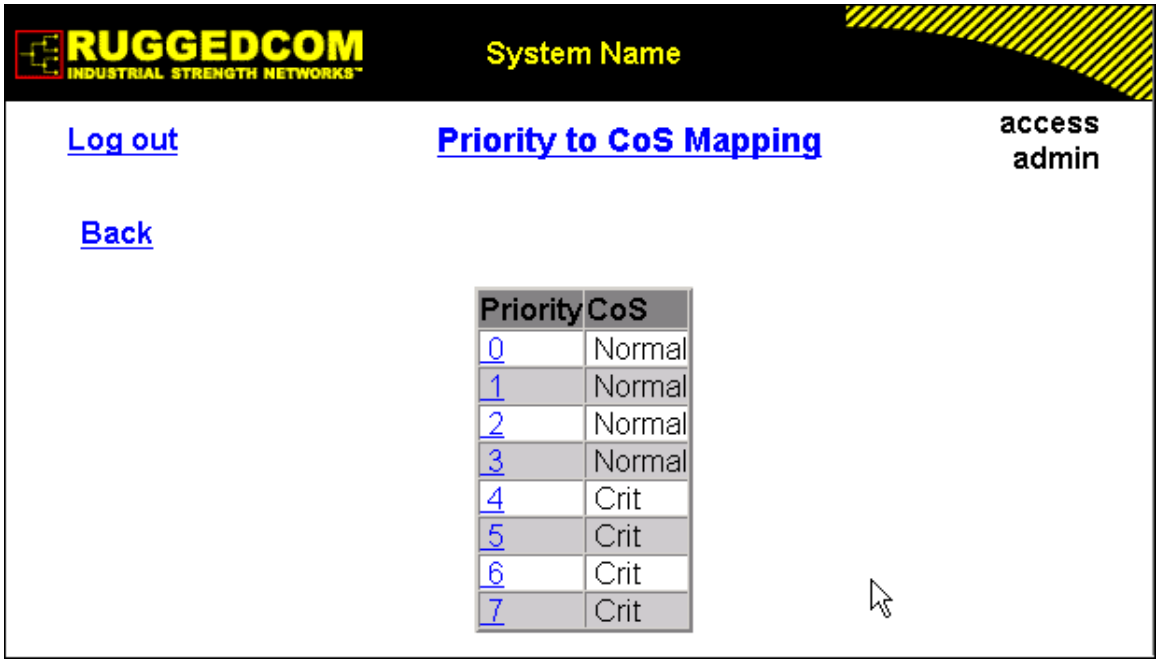

Figure 8.6. Priority to CoS Mapping Table

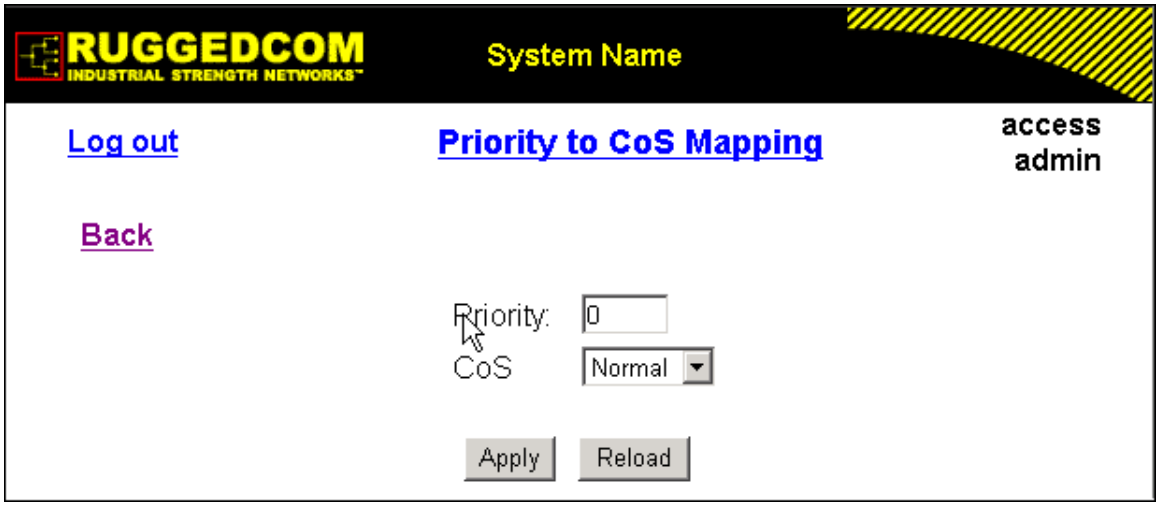

Figure 8.7. Priority to CoS Mapping Form

Priority

```
Synopsis: 0 to 7
Default: 0
```

This is a value of the IEEE 802.1p priority.

CoS

```
Synopsis: { Normal, Medium, High, Crit }
Default: Normal
```

This is a CoS assigned to received tagged frames with the specified IEEE 802.1p priority value.

8.2.4. DSCP to CoS Mapping

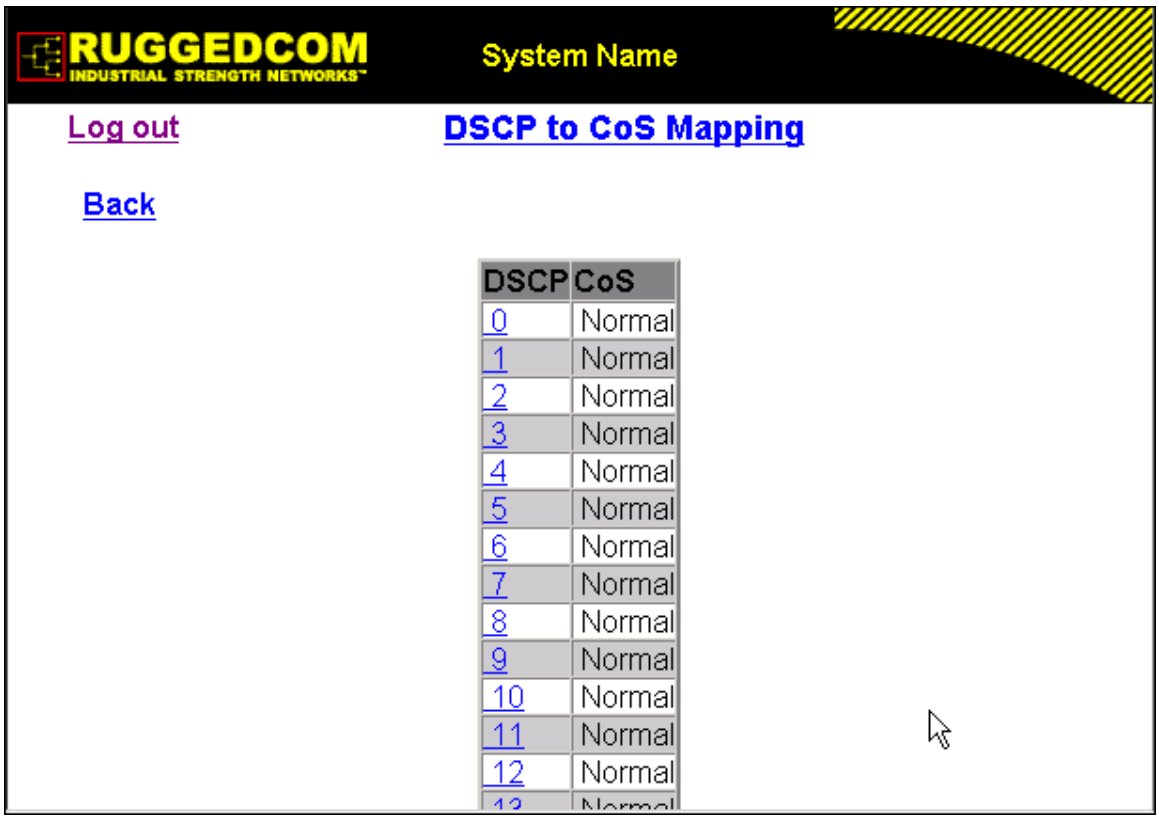

Figure 8.8. TOS DSCP to CoS Mapping Table

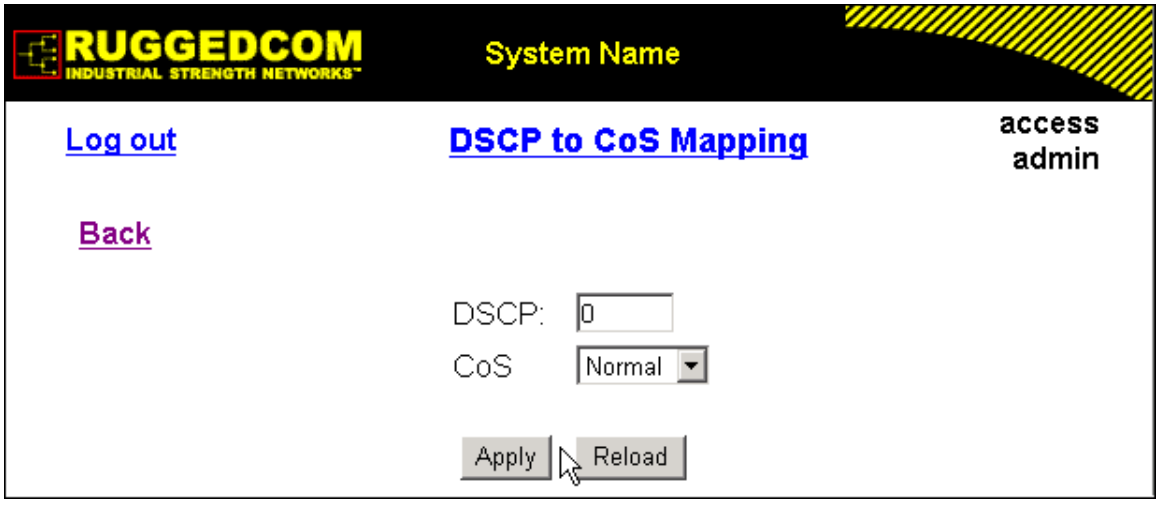

Figure 8.9. TOS DSCP to CoS Mapping Form

DSCP

```
Synopsis: 0 to 63
Default: 0
```

This is a Differentiated Services Code Point (DSCP) - a value of the 6-bit DiffServ field in the Type-Of-Service (TOS) field of the IP header.

CoS

```
Synopsis: { Normal, Medium, High, Crit }
Default: Normal
```

This is a Class of Service assigned to received frames with the specified DSCP.

9. Multicast Filtering

ROS® Multicast Filtering provides the following features:

- Support for up to 256 Multicast Groups (either static or dynamic).
- Ability to prioritize a Static Multicast Group via Class-of-Service.
- Industry standard support of IGMP (RFC 1112, RFC 2236) versions 1 and 2 in active and passive roles.
- Support of IEEE 802.1Q-2005 standard GMRP (GARP Multicast Registration protocol).
- Ability to enable or disable IGMP on a per VLAN basis.
- Multicast routers may be statically configured or dynamically recognized by IGMP.
- “Routerless” IGMP operation.

ROS performs Multicast Filtering using the following methods:

- Static Multicast Groups.
- Internet Group Management Protocol (IGMP) snooping.
- IEEE standard GARP Multicast Registration protocol (GMRP).

Note

ROS IGMP Snooping supports multicast routers using IGMP version 2 and hosts using either IGMP version 1 or 2.

9.1. IGMP

IGMP is used by IP hosts to report their host group memberships to multicast routers. As hosts join and leave specific multicast groups, streams of traffic are directed to or withheld from that host.

The IGMP protocol operates between multicast routers and IP hosts. When an unmanaged switch is placed between multicast routers and their hosts, the multicast streams will be distributed to all ports. This may introduce significant traffic onto ports that do not require it and receive no benefit from it.

RuggedCom products with IGMP Snooping enabled will act upon IGMP messages sent from the router and the host, restricting traffic streams to the appropriate LAN segments.

9.1.1. Router and Host IGMP Operation

The network shown in Figure 9.1, “IGMP Operation Example 1” provides a simple example of the use of IGMP. One “producer” IP host (P1) is generating two IP multicast streams, M1 and M2. There are four potential “consumers” of these streams, C1 through C4.

The multicast router discovers which host wishes to subscribe to which stream by sending general membership queries to each of the segments.

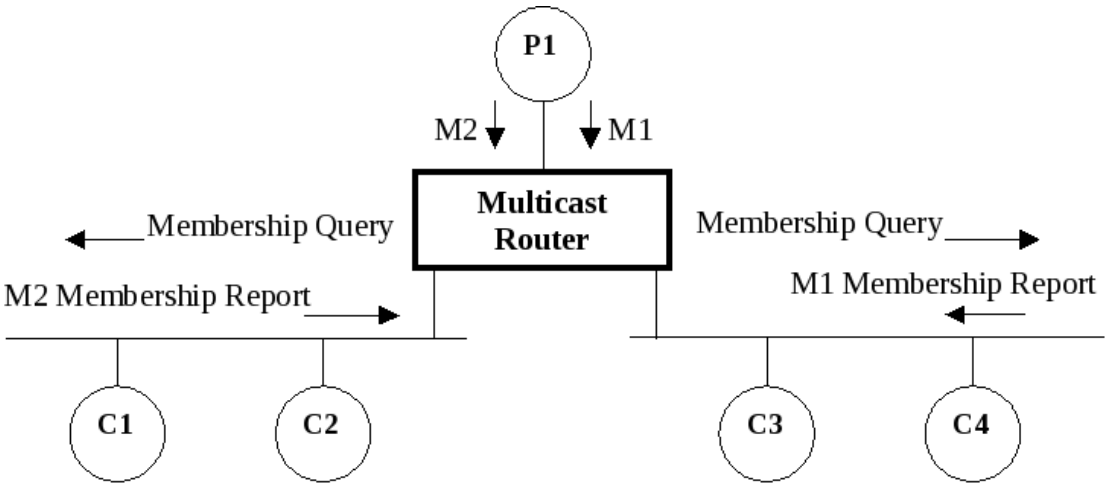

Figure 9.1. IGMP Operation Example 1

In this example, the general membership query sent to the C1-C2 segment is answered by a membership report indicating the desire to subscribe to a stream M2. The router will forward the M2 stream onto the C1-C2 segment. In a similar fashion, the router discovers that it must forward M1 onto segment C3-C4.

Note

Membership reports are also referred to as "joins".

A consumer may join any number of multicast groups, issuing a membership report for each group. When a host issues a membership report, other hosts on the same network segment that also require membership to the same group suppress their own requests, since they would be redundant. In this way, the IGMP protocol guarantees that the segment will issue only one join for each group.

The router periodically queries each of its segments in order to determine whether at least one consumer still subscribes to a given stream. If it receives no responses within a given timeout period (usually two query intervals), the router will prune the multicast stream from the given segment.

A more usual method of pruning occurs when consumers wishing to unsubscribe issue an IGMP "leave group" message. The router will immediately issue a group-specific membership query to determine whether there are any remaining subscribers of that group on the segment. After the last consumer of a group has un-subscribed, the router will prune the multicast stream from the given segment.

9.1.2. Switch IGMP Operation

The IGMP Snooping feature provides a means for switches to snoop (i.e. watch) the operation of routers, respond with joins/leaves on the behalf of consumer ports and to prune multicast streams accordingly.

There are two modes of IGMP that the switch can be configured to assume - active and passive.

Active Mode

ROS IGMP supports “routerless” mode of operation.

When such a switch is used without a multicast router, it is able to function as if it is a multicast router sending IGMP general queries.

Passive Mode

When such a switch is used in a network with a multicast router, it can be configured to run Passive IGMP. This mode prevents the switch from sending the queries that can confuse the router causing it to stop issuing IGMP queries.

Note

A switch running in passive mode requires the presence of a multicast router or it will not be able to forward multicast streams at all

If no multicast routers are present, at least one IGMP Snooping switch must be configured for Active IGMP mode to make IGMP functional.

IGMP Snooping Rules

- When a multicast source starts multicasting, the traffic stream will be immediately blocked on segments from which joins have not been received.
- The switch will always forward all multicast traffic to the ports where multicast routers are attached unless configured otherwise.
- Packets with a destination IP multicast address in the 224.0.0.X range which are not IGMP are always forwarded to all ports. This behavior is based on the fact that many systems do not send joins for IP multicast addresses in this range while still listening to such packets.
- The switch implements “proxy-reporting”, i.e. membership reports received from downstream are summarized and used by the switch to issue its own reports.
- The switch will only send IGMP membership reports out of those ports where multicast routers are attached because sending membership reports to hosts could result in unintentionally preventing a host from joining a specific group.
- Multicast routers use IGMP to elect a master router known as the querier – the one with the lowest IP address is elected to be the querier, all other routers become of non-queriers, participating only forward multicast traffic. Switches running in Active IGMP mode participate in the querier election like multicast routers.
- When the querier election process is complete, the switch simply relays IGMP queries received from the querier.
- When sending IGMP packets, the switch uses its own IP address, if it has one, for the VLAN on which packets are sent, or an address of 0.0.0.0, if it doesn’t have an assigned IP address.

Note

IGMP Snooping switches perform multicast pruning using a multicast frames’ destination MAC multicast address which depends on the group IP multicast address. IP address W.X.Y.Z

corresponds to MAC address 01-00-5E-XX-YY-XX where XX is the lower 7 bits of X and YY and ZZ are simply Y and Z coded in hexadecimal.

One can note that IP multicast addresses such as 224.1.1.1 and 225.1.1.1 will both map onto the same MAC address 01-00-5E-01-01-01. This is indeed a problem for which the IETF Network Working Group currently has offered no solution. Users are advised to be aware of and avoid this problem.

IGMP and RSTP

An RSTP change of topology can render the routes selected to carry multicast traffic as incorrect. This results in lost multicast traffic.

If RSTP detects change in the network topology, IGMP will take some actions to avoid loss of multicast connectivity and reduce network convergence time:

- The switch will immediately issue IGMP queries (if in IGMP Active mode) to obtain potential new group membership information.
- The switch can be configured to flood multicast streams temporarily out of all ports that are not configured as RSTP Edge Ports.

9.1.3. Combined Router and Switch IGMP Operation

This section describes the additional challenges of multiple routers, VLAN support and switching.

Producer P1 resides upon VLAN 2 while P2 resides upon VLAN 3. Consumer C1 resides upon both VLANs whereas C2 and C3 reside upon VLANs 3 and 2, respectively. Router 2 resides upon VLAN 2, presumably to forward multicast traffic to a remote network or act as a source of multicast traffic itself.

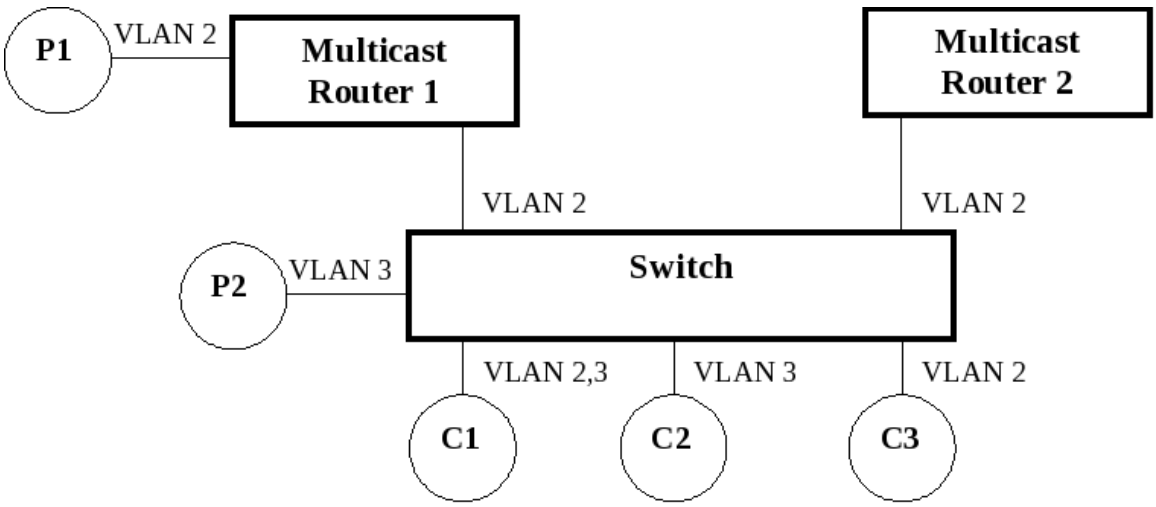

Figure 9.2. IGMP Operation Example 2

In this example, we will assume that all the devices agree that router 1 is the querier for VLAN 2 and router 2 is simply a non-querier. In this case, the switch will periodically receive queries from router 1 and, thus, maintain the information concerning which of its ports links to the multicast

router. However, the switch port that links to router 2 must be manually configured as a “router port”. Otherwise, the switch will send neither multicast streams nor joins/leaves to router 2.

Note that VLAN 3 does not have an external multicast router. The switch should be configured to operate in its “routerless” mode and issue general membership queries as if it is the router.

Processing Joins

If host C1 desires to subscribe to the multicast streams for both P1 and P2, it will generate two joins. The join from C1 on VLAN 2 will cause the switch to immediately initiate its own join to multicast router 1 (and to issue its own join as a response to queries).

The join from C1 for VLAN 3 will cause the switch to immediately begin forwarding multicast traffic from P2 to C2.

Processing Leaves

When host C1 decides to leave a multicast group, it will issue a leave request to the switch. The switch will poll the port to determine if C1 is the last member of the group on that port. If C1 is the last (or only) member, the group will immediately be pruned from the port.

Should host C1 leave the multicast group without issuing a leave group message and then fail to respond to a general membership query, the switch will stop forwarding traffic after two queries.

When the last port in a multicast group leaves the group (or is aged-out), the switch will issue an IGMP leave report to the router.

9.2. GMRP (GARP Multicast Registration Protocol)

The GARP Multicast Registration Protocol (GMRP) is an application of the Generic Attribute Registration Protocol (GARP) that provides a mechanism at Layer 2 for managing multicast group membership in a bridged Layer 2 network. It allows Ethernet switches and end stations to register and unregister membership in multicast groups with other switches on a LAN, and for that information to be disseminated to all switches in the LAN that support Extended Filtering Services.

GMRP is an industry-standard protocol first defined in IEEE 802.1D-1998 and extended in IEEE 802.1Q-2005. GARP was defined in IEEE 802.1D-1998 and updated in 802.1D-2004. Note that GMRP provides similar functionality at Layer 2 to that which IGMP, described in the preceding sections, provides at Layer 3.

Joining a Multicast Group)

In order to join a multicast group, an end station transmits a GMRP “join” message. The switch that receives the “join” message adds the port through which the message was received to the multicast group specified in the message. It then propagates the “join” message to all other hosts in the VLAN, one of which is expected to be the multicast source.

When a switch transmits GMRP updates (from GMRP-enabled ports), all of the multicast groups known to the switch, whether configured manually or learned dynamically through GMRP, are advertised to the rest of network.

As long as one host on the Layer 2 network has registered for a given multicast group, traffic from the corresponding multicast source will be carried on the network. Traffic multicast by the source is only forwarded by each switch in the network to those ports from which it has received join messages for the multicast group.

Leaving a Multicast Group

Periodically, the switch sends GMRP queries in the form of a “leave all” message. If a host (either a switch or an end station) wishes to remain in a multicast group, it reasserts its group membership by responding with an appropriate “join” request. Otherwise, it can either respond with a “leave” message or simply not respond at all. If the switch receives a “leave” message or receives no response from the host for a timeout period, the switch removes the host from the multicast group.

GMRP Protocol Notes

Since GMRP is an application of GARP, transactions take place using the GARP protocol. GMRP defines the following two Attribute Types:

- The **Group Attribute Type**, used to identify the values of group MAC addresses
- The **Service Requirement Attribute Type**, used to identify service requirements for the group

Service Requirement Attributes are used to change the receiving port’s multicast filtering behavior to one of the following:

- Forward All Multicast group traffic in the VLAN, or
- Forward All Unknown Traffic (Multicast Groups) for which there are no members registered in the device in a VLAN

If GMRP is disabled on the RuggedSwitch® then GMRP PDUs received by the switch will be forwarded like any other traffic; but if GMRP is enabled on at least one of the ports, then GMRP packets will be processed by the switch, and not forwarded.

9.2.1. GMRP Example

In the example depicted in Figure 9.3, “Example using GMRP”, there are two multicast sources, S1 and S2, multicasting to Multicast Groups 1 and 2, respectively. A network of five switches, including one core Switch, B, connects the sources to two hosts, H1 and H2, which receive the multicast streams from S1 and S2, respectively.

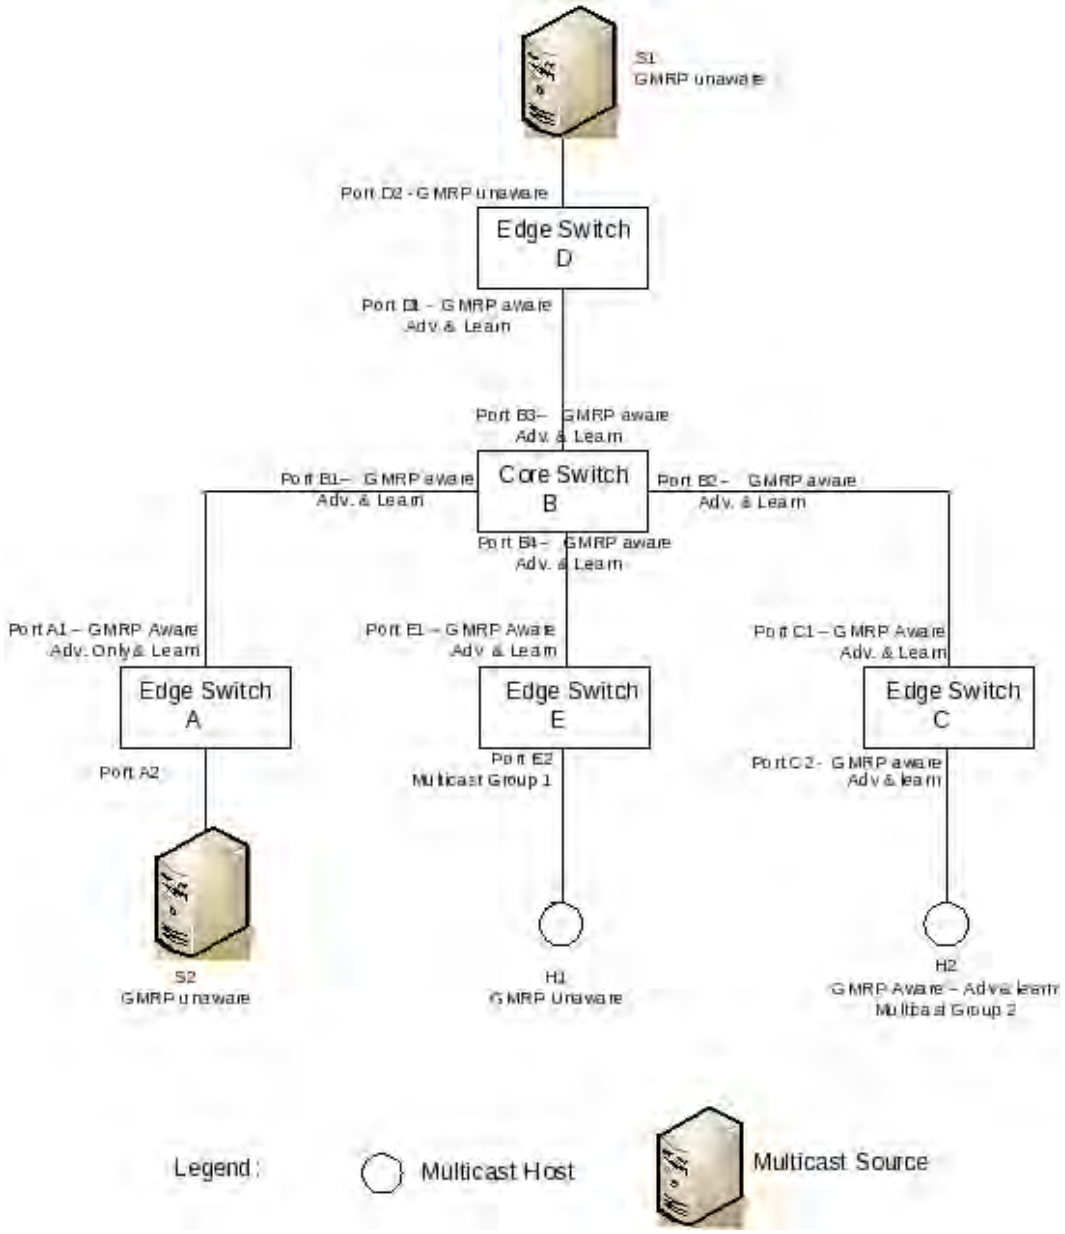

Figure 9.3. Example using GMRP

Joining the Multicast Groups:

The sequence of events surrounding the establishment of membership for the two Multicast Groups on the example network is as follows:

- Host H1 is GMRP unaware but needs to see traffic for Multicast Group 1. Port E2 on Switch E, therefore, is statically configured to forward traffic for Multicast Group 1.
- Switch E advertises membership in Multicast Group 1 to the network through Port E1, making Port B4 on Switch B a member of Multicast Group 1.
- Switch B propagates the “join” message, causing Port D1 on Switch D to become a member of Multicast Group 1. Note that ports A1 and C1 also become members.

- Host H2 is GMRP-aware and sends a “join” request for Multicast Group 2 to Port C2, which thereby becomes a member of Group 2.
- Switch C propagates the “join” message, causing Port B2 on Switch B and Port A1 on Switch A to become members of Multicast Group 2. Note that ports D1 and E1 also become members.

Multicast Traffic on the Network

Once GMRP-based registration has propagated through the network as described above, multicasts from S1 and S2 can reach their destinations, as described in the following:

- Source S1 transmits multicast traffic to Port D2 which is forwarded via Port D1, which has previously become a member of Multicast Group 1.
- Switch B forwards the Group 1 multicast via Port B4 towards Switch E.
- Switch E forwards the Group 1 multicast via Port E2, which has been statically configured for membership in Multicast Group 1.
- Host H1, connected to Port E2, thus receives the Group 1 multicast.
- Source S2 transmits multicast traffic to Port A2, which is then forwarded via port A1, which has previously become a member of Multicast Group 2.
- Switch B forwards the Group 2 multicast via Port B2 towards Switch C.
- Switch C forwards the Group 2 multicast via Port C2, which has previously become a member of Group 2.
- Ultimately, Host H2, connected to Port C2, receives the Group 2 multicast.

9.3. Multicast Filtering Configuration and Status

The Multicast Filtering menu is available from the main menu.

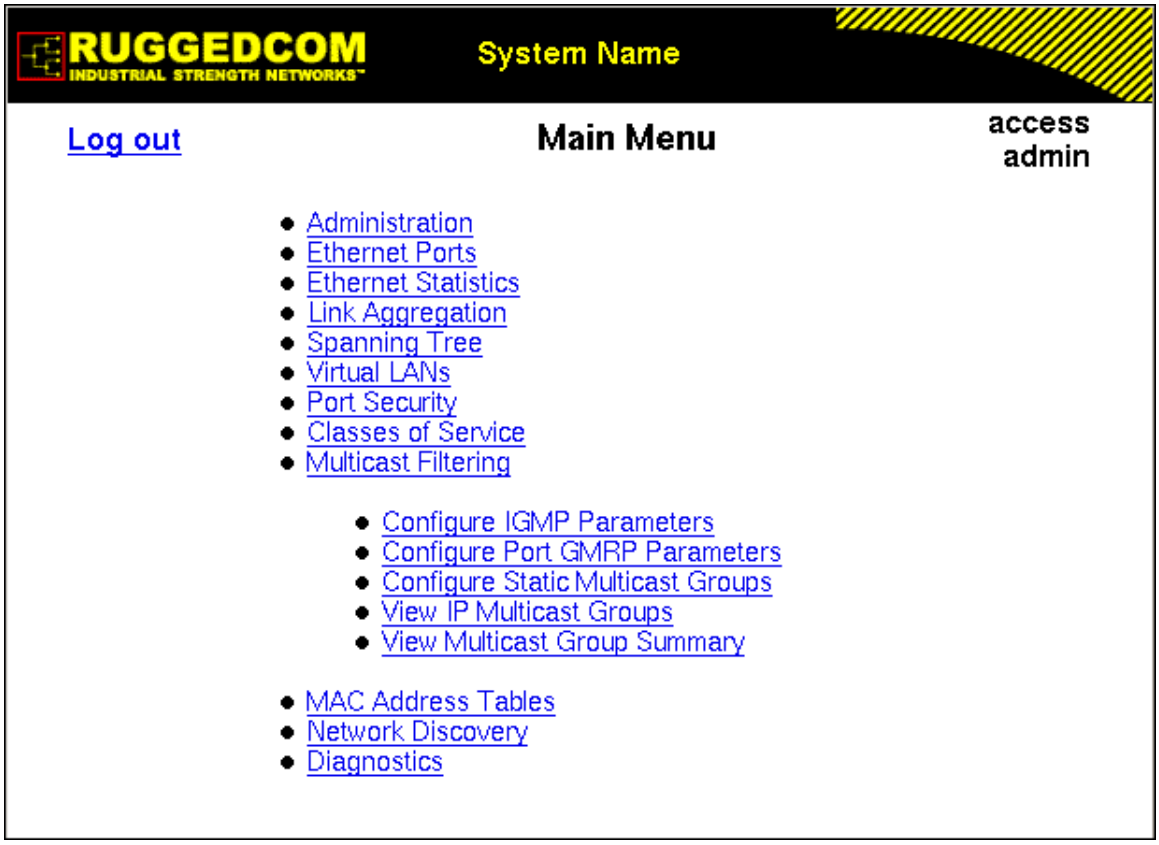

Figure 9.4. Multicast Filtering Menu

9.3.1. Configuring IGMP Parameters

Note that the activation of IGMP on a per-VLAN basis is configured using Static VLANs.

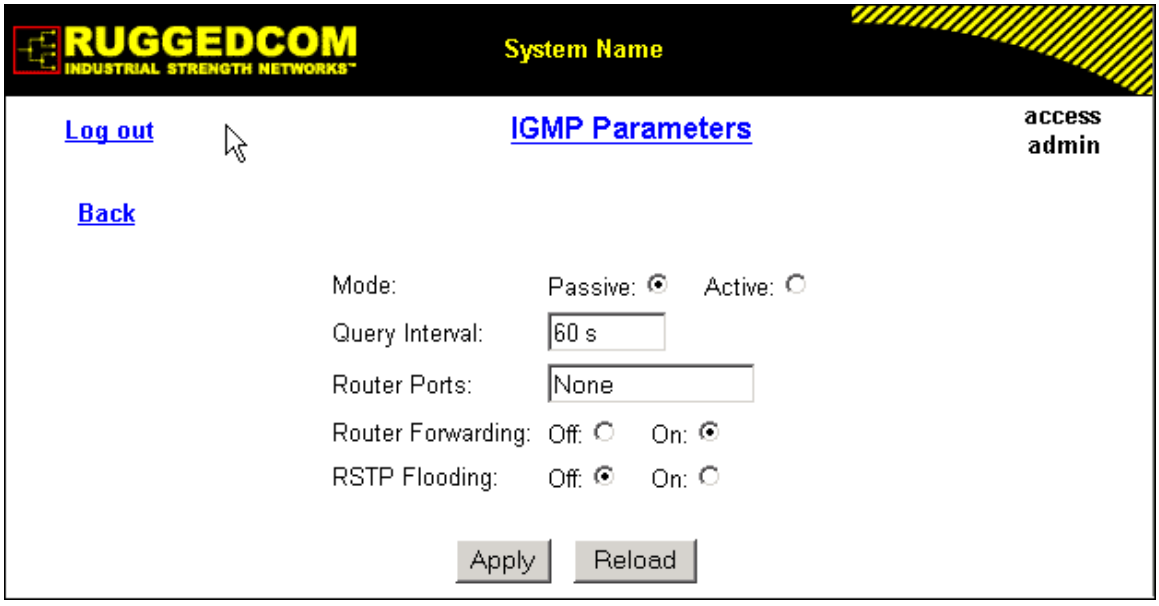

Figure 9.5. IGMP Parameters Form

Mode

```
Synopsis: { Passive, Active }
Default: Passive
```

Specifies IGMP mode:

PASSIVE - the switch passively snoops IGMP traffic and never sends IGMP queries

ACTIVE - the switch generates IGMP queries, if no queries from a better candidate for being the querier are detected for a while.

Query Interval

```
Synopsis: 10 to 3600
Default: 60 s
```

The time interval between IGMP queries generated by the switch.

Note

This parameter also affects the Group Membership Interval (i.e. the group subscriber aging time), therefore, it takes effect even in PASSIVE mode.

Router Ports

```
Synopsis: Any combination of numbers valid for this parameter
Default: None
```

This parameter specifies ports that connect to multicast routers. If you do not configure known router ports, the switch may be able to detect them, however it is advisable to pre-configure them.

Router Forwarding

```
Synopsis: { Off, On }
Default: On
```

This parameter specifies whether multicast streams will be always forwarded to multicast routers.

RSTP Flooding

```
Synopsis: { Off, On }
Default: Off
```

This parameter specifies whether multicast streams will be flooded out of all RSTP non-edge ports upon topology change detection. Such flooding is desirable, if guaranteed multicast stream delivery after topology change is most important.

9.3.2. Configuring Static Multicast Groups

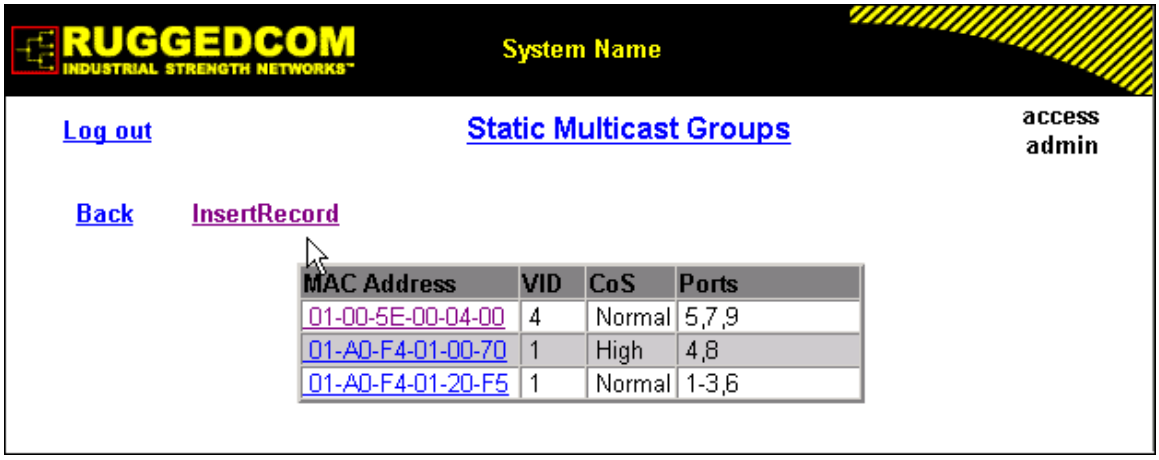

Figure 9.6. Static Multicast Groups Table

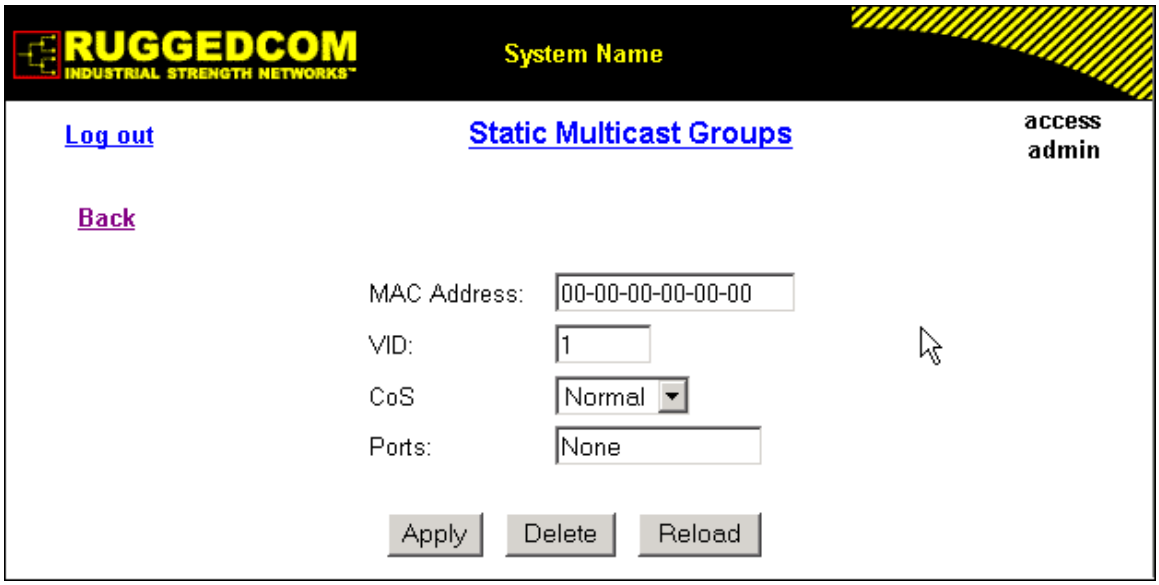

Figure 9.7. Static Multicast Group Form

MAC Address

```
Synopsis: ##-##-##-##-##-## where ## ranges 0 to FF
Default: 00-00-00-00-00-00
```

A multicast group MAC address.

VID

```
Synopsis: 1 to 4094
Default: 1
```

The VLAN Identifier of the VLAN upon which the multicast group operates.

CoS

```
Synopsis: { Normal, Medium, High, Crit }
```

`Default: Normal`

Specifies what Class Of Service is assigned to the multicast group frames.

Ports

`Synopsis: Any combination of numbers valid for this parameter`
`Default: None`

The ports to which the multicast group traffic is forwarded.

9.3.3. Viewing IP Multicast Groups

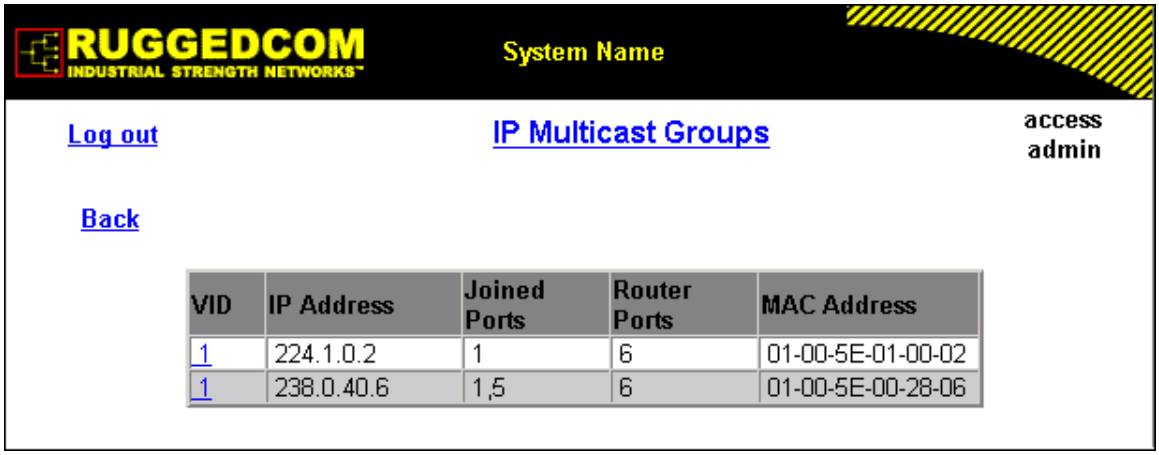

Figure 9.8. IP Multicast Groups Table

VID

`Synopsis: 0 to 65535`

The VLAN Identifier of the VLAN upon which the multicast group operates.

IP Address

`Synopsis: ###.###.###.### where ### ranges from 0 to 255`

The multicast group IP address.

Joined Ports

`Synopsis: Any combination of numbers valid for this parameter`

All ports that subscribed to the multicast group traffic.

Router Ports

`Synopsis: Any combination of numbers valid for this parameter`

All ports that have been manually configured or dynamically discovered (by observing router specific traffic) as ports that link to multicast routers.

MAC Address

`Synopsis: ##-##-##-##-##-## where ## ranges 0 to FF`

The multicast MAC address corresponding to the group multicast IP address.

9.3.4. Configuring GMRP

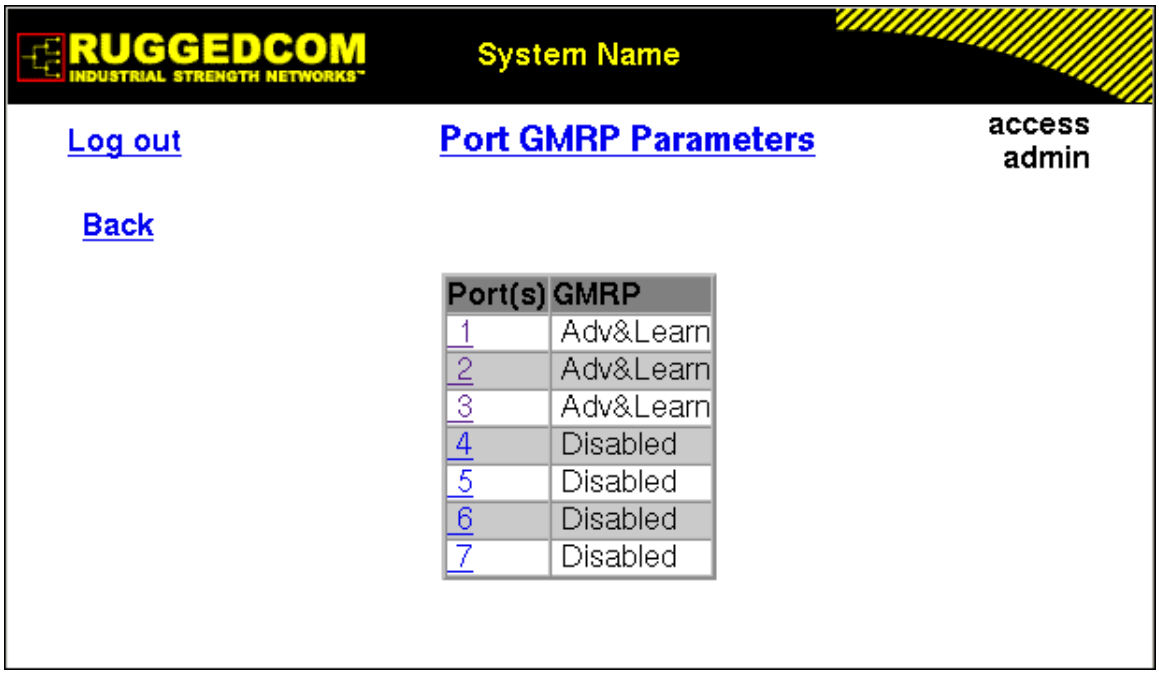

Figure 9.9. GMRP Port Summary

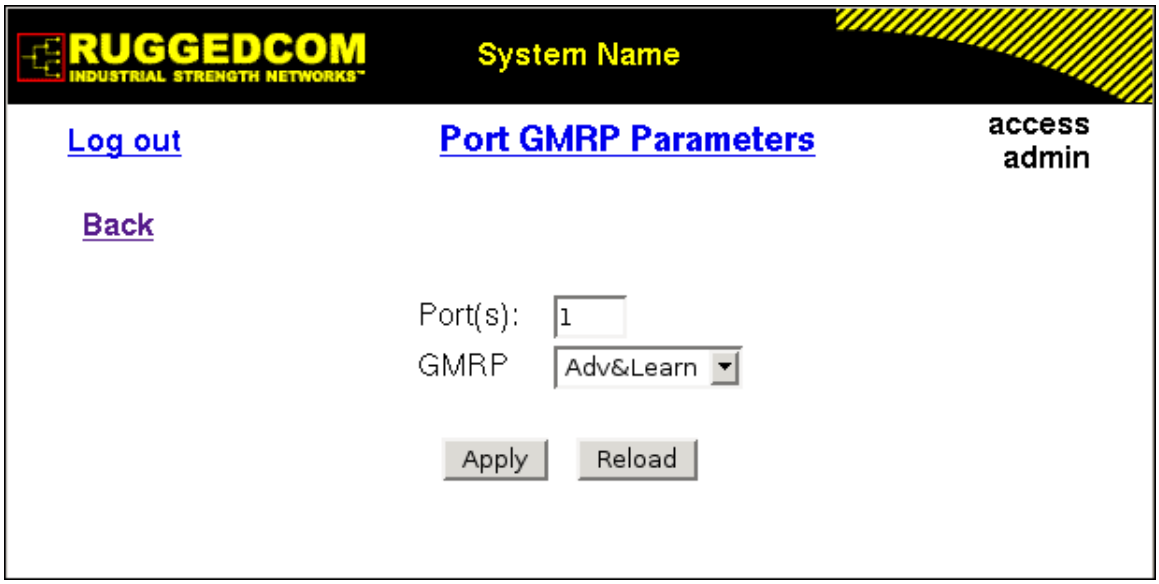

Figure 9.10. Port GMRP Parameters

Port(s)

`Synopsis: Any combination of numbers valid for this parameter`

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

GMRP

```
Synopsis: { Disabled, Adv Only, Adv&Learn }
Default: Disabled
```

Configures GMRP (GARP Multicast Registration Protocol) operation on the port. There are three GMRP modes of operation:

DISABLED - the port is not capable of any GMRP processing.

ADVERTISE ONLY - the port will declare all MCAST addresses existing in the switch (configured or learned) but will not learn any MCAST addresses.

ADVERTISE & LEARN - the port will declare all MCAST Addresses existing in the switch (configured or learned) and can dynamically learn MCAST addresses.

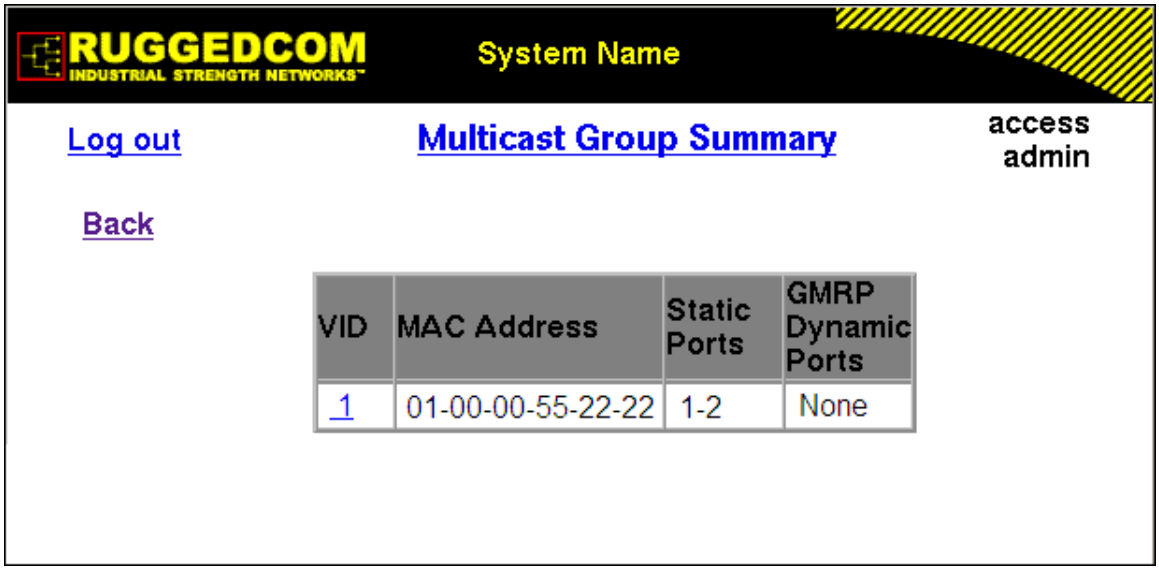

Figure 9.11. Multicast Group Summary

VID

`Synopsis: 0 to 65535`

The VLAN Identifier of the VLAN upon which the multicast group operates.

MAC Address

`Synopsis: ##-##-##-##-##-## where ## ranges 0 to FF`

The multicast group MAC address.

Static Ports

`Synopsis: Any combination of numbers valid for this parameter`

Ports that joined this group statically through static configuration in Static MAC Table and to which the multicast group traffic is forwarded.

GMRP Dynamic Ports

`Synopsis: Any combination of numbers valid for this parameter`

Ports that joined this group dynamically through GMRP Application and to which the multicast group traffic is forwarded.

9.4. Troubleshooting

Problem One

When I start a multicast traffic feed, it is always distributed to all members of the VLAN.

Is IGMP enabled for the VLAN? Multicasts will be distributed to all members of the VLAN unless IGMP is enabled.

Problem Two

Computers on my switch receive the multicast traffic just fine, but I can't get the stream through a connected router.

Is the port used to connect the router included in the Router Ports list?

To determine whether the multicast stream is being delivered to the router, run the Ethernet Statistics menu **View Ethernet Statistics** command. Verify that the traffic count transmitted to the router is the same as the traffic count received from the multicasting source.

Problem Three

The video stream at one of my end stations is of pretty poor quality.

Video serving is a resource-intensive application. Because it uses isochronous workload, data must be fed at a prescribed rate or end users will see glitches in the video. Networks that carry data from the server to the client must be engineered to handle this heavy, isochronous workload.

Video streams can consume large amounts of bandwidth. Features and capacity of both server and network (including routers, bridges, switches, and interfaces) impact the streams.

You should not exceed 60% of the maximum interface bandwidth. For example, if using a 10 Mbps Ethernet, you should run a single multicasting source at no more than 6 Mbps, or two sources at 3 Mbps.

Router ports will carry the traffic of all multicast groups, so it is especially important to consider these ports in your design.

Note that multicasting will definitely introduce latency in all traffic on the network. Plan your network carefully in order to account for capacity and latency concerns.

Problem Four

Multicast streams of some groups are not forwarded properly. Some segments without subscribers receive the traffic while some segments with subscribers don't.

Ensure that you do not have a situation where different multicast groups have multicast IP addresses that map to the same multicast MAC address. The switch forwarding operation is MAC address-based and will not work properly for several groups mapping to the same MAC address.

Problem Five

Computers on my switch issue join requests but don't receive multicast streams from a router.

Is your multicast router running IGMP version 2? It must run IGMP version 2 in order for IGMP Snooping to operate properly.

Problem Six

I connect or disconnect some switch ports and multicast goes everywhere. Is IGMP broken?

No, it may be a proper switch behavior. When the switch detects a change in the network topology through RSTP, it acts to avoid loss of multicast traffic – if configured to do so, it starts forwarding all multicast traffic to all ports that are not RSTP Edge ports (because they may potentially link to routers). This may result in some undesired flooding of multicast traffic (which will stop after a few minutes), however, it guarantees that all devices interested in the traffic will keep receiving it without a break. Note that the same behavior will be observed when the switch resets or when IGMP Snooping is being enabled for the VLAN.

10. MAC Address Tables

ROS® MAC address table management provides you with the following features:

- Viewing learned MAC addresses.
- Purging MAC Address Entries.
- Configuring the switch's MAC Address Aging time.
- Configuring static MAC addresses.

The MAC Address Tables menu is accessible from the main menu.

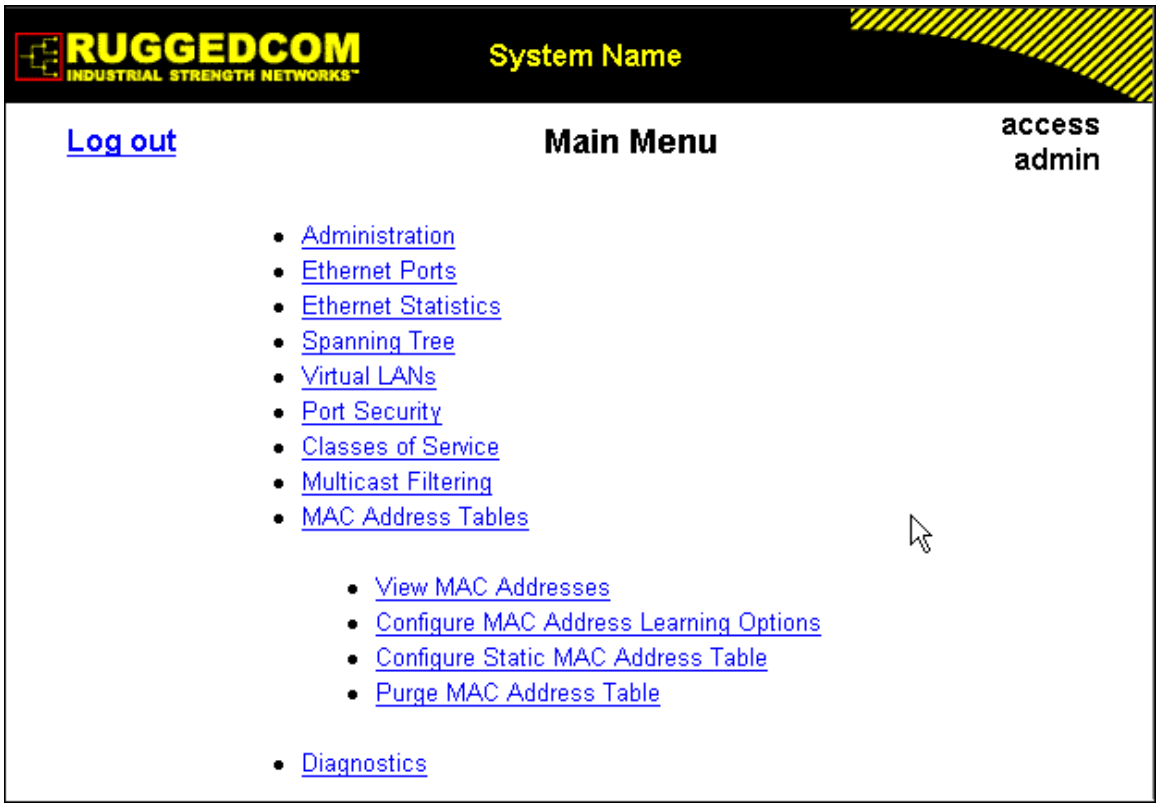

Figure 10.1. MAC Address Tables Menu

10.1. Viewing MAC Addresses

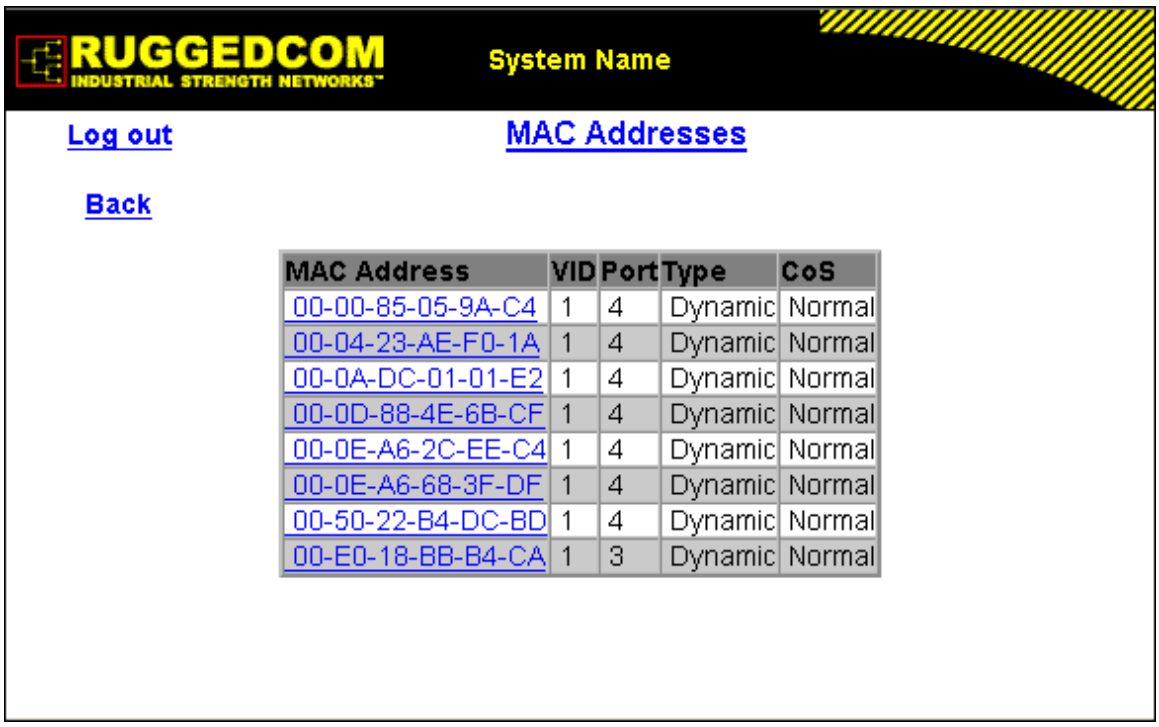

Figure 10.2. Address Table

MAC Address

Synopsis: ##-##-##-##-##-## where ## ranges 0 to FF

A MAC address learned by the switch.

VID

Synopsis: 0 to 65535

The VLAN Identifier of the VLAN upon which the MAC address operates.

Port

Synopsis: 0 to 65535 or { Multi, Local }

The port on which MAC address has been learned.

MULTI - multicast address, so there is no switch port associated with this MAC address.

Type

Synopsis: { Static, Dynamic }

This describes how the MAC address has been learned by the switch:

STATIC - the address has been learned as a result of a Static MAC Address Table configuration or some other management activity and can not be automatically unlearned or relearned by the switch.

DYNAMIC - The address has been automatically learned by the switch and can be automatically unlearned.

CoS

Synopsis: { Normal, Medium, High, Crit }

Specifies what Class Of Service is assigned to frames carrying this address as source or destination address.

10.2. Configuring MAC Address Learning Options

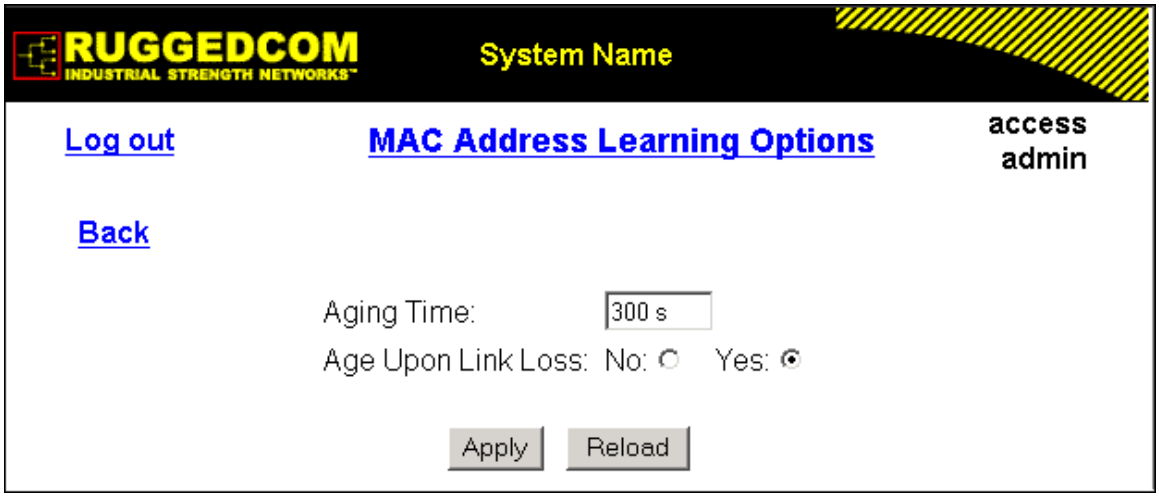

Figure 10.3. MAC Address Learning Options Form

Aging Time

```
Synopsis: 15 to 800
Default: 300 s
```

This parameter configures the time that a learned MAC address is held before being aged out.

Age Upon Link Loss

```
Synopsis: { No, Yes }
Default: Yes
```

When a link failure (and potentially a topology change) occurs, the switch may have some MAC addresses previously learned on the failed port. As long as those addresses are not aged out, the switch will still be forwarding traffic to that port, thus preventing that traffic from reaching its destination via the new network topology. This parameter allows the aging out of all MAC addresses learned on a failed port immediately upon link failure detection.

10.3. Configuring Static MAC Address Table

Static MAC addresses are usually configured when the user wishes to enforce port security (if supported).

Static MAC addresses are also configured when a device can receive but cannot transmit frames.

Prioritized MAC addresses are configured when traffic to or from a specific device on a LAN segment is to be assigned a higher CoS priority than other devices on that LAN segment.

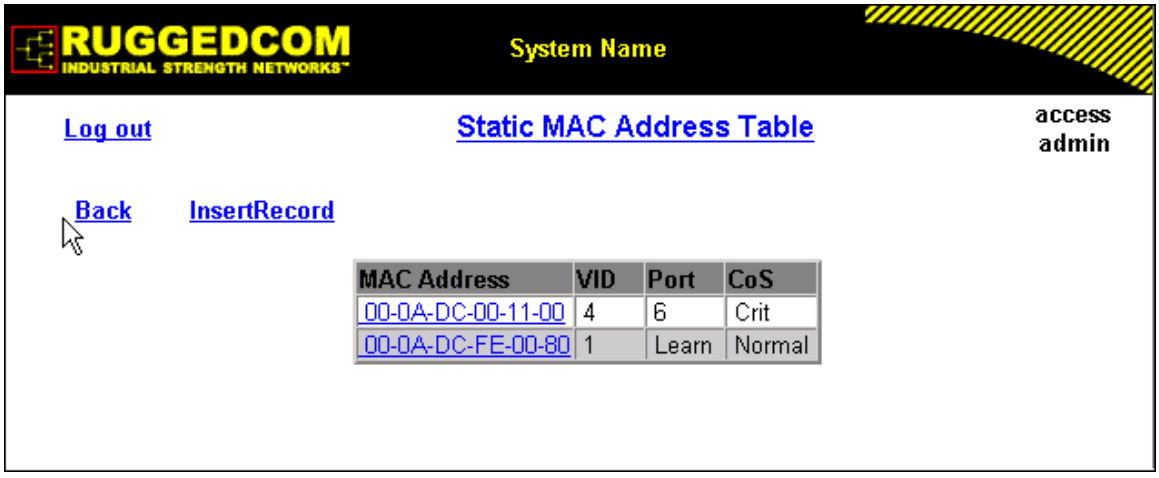

Figure 10.4. Static MAC Address Table

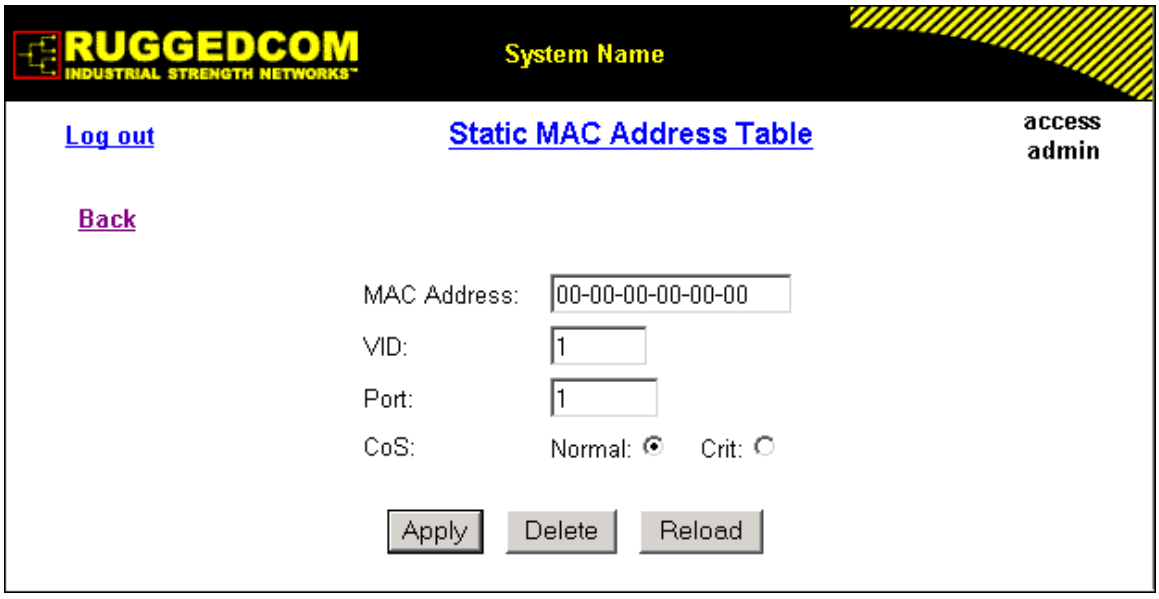

Figure 10.5. Static MAC Address Form

MAC Address

```
Synopsis: ##-##-##-##-##-## where ## ranges 0 to FF
Default: 00-00-00-00-00-00
```

A MAC address that is to be statically configured.

VID

```
Synopsis: 1 to 1000
Default: 1
```

The VLAN Identifier of the VLAN upon which the MAC address operates.

Port

```
Synopsis: 1 to maximum port number or { Learn }
Default: Learn
```

Enter the port number upon which the device with this address is located. If the port should be auto-learned, set this parameter to 'Learn'.

CoS

```
Synopsis: { Normal, Medium, High, Crit }
Default: Normal
```

Set this parameter to prioritize the traffic for a specified address.

10.4. Purging MAC Address Table

This command removes all dynamic entries from the MAC address table. The only negative impact of this operation is that it causes flooding while addresses are relearned.

11. Network Discovery

ROS® supports two different Layer 2 protocols for automated network discovery: LLDP, (the Link Layer Discovery Protocol) and RCDP (the RuggedCom Discovery Protocol™).

LLDP is an IEEE standard protocol, IEEE 802.11AB, which allows a networked device to advertise its own basic networking capabilities and configuration. ROS is capable of advertising and collecting network information via LLDP. LLDP functionality in ROS includes the ability to:

- Enable or disable LLDP reception and transmission per port or for the whole device.
- View LLDP statistics.
- View 'neighbor' information.
- Report LLDP neighbor information via SNMP.

RCDP (the RuggedCom Discovery Protocol) is designed primarily for the initial deployment of (unconfigured) RuggedCom networking devices. In response to RCDP commands and queries from an application such as RuggedExplorer™, which supports RCDP, ROS has the ability to:

- Enable or disable RCDP functionality.
- Report its basic network configuration and other identifying information.
- Respond to a basic set of control commands.
- Perform basic device configuration.

11.1. LLDP Operation

The IEEE standard, 802.1AB Link Layer Discovery Protocol (LLDP), describes a protocol that can simplify the troubleshooting of complex networks and can be used by Network Management Systems (NMS) to obtain and monitor detailed information about a network's topology. LLDP data are made available via SNMP (through support of LLDP-MIB).

LLDP is a neighbor discovery protocol. It defines a standard method for Ethernet network devices such as switches and routers to advertise information about themselves to other nodes on the network and to store the information they discover. Details such as device configuration, device capabilities and device identification can be advertised using this protocol.

LLDP agent operation is typically implemented as two modules: the LLDP transmit module and LLDP receive module. The LLDP transmit module, when enabled, sends the local device's information at regular intervals, in 802.1AB standard format. Whenever the transmit module is disabled, it transmits an LLDPDU (LLDP data unit) with a time-to-live (TTL) TLV containing "0" in the information field. This enables remote devices to remove the information associated with the local device in their databases. The LLDP receive module, when enabled, receives remote devices' information and updates its LLDP database of remote systems. When new or updated information is received, the receive module initiates a timer for the valid duration indicated by the TTL TLV in the received LLDPDU. A remote system's information is removed from the database when an LLDPDU is received from it with TTL TLV containing "0" in its information field.

Note

LLDP is implemented to keep a record of only one device per Ethernet port. Therefore, if there are multiple devices sending LLDP information to a switch port on which LLDP is enabled, information about the neighbor on that port will change constantly.

11.2. RCDP Operation

The purpose of the RuggedCom Discovery Protocol™ is to support the deployment of ROS-based devices that have not been configured since leaving the factory. Unconfigured ROS devices all have the default IP (Layer 3) address. Connecting more than one of them on a Layer 2 network means that one cannot use standard IP-based configuration tools to configure them. The behavior of IP-based mechanisms such as the web interface, SSH, telnet, or SNMP will all be undefined.

Since RCDP operates at Layer 2, it can be used to reliably and unambiguously address multiple devices even though they may share the same IP configuration.

RuggedCom's RuggedExplorer™ is a lightweight, standalone Windows application that supports RCDP. It is capable of discovering, identifying and performing basic configuration of ROS-based devices via RCDP. The features supported by RCDP include:

- Discovery of ROS-based devices over a Layer 2 network.
- Retrieval of basic network configuration, ROS version, order code, and serial number.
- Control of device LEDs for easy physical identification.
- Configuration of basic identification, networking, and authentication parameters.

Note

RCDP is not compatible with VLAN-based network configurations. For correct operation of RuggedExplorer, no VLANs (tagged or untagged) must be configured. All VLAN configuration items must be at their default settings.

Note

ROS responds to RCDP requests only - it does not under any circumstances initiate any RCDP-based communication.

11.3. Network Discovery Menu

The main Network Discovery menu links to configuration menus for both LLDP and RCDP.

Figure 11.1. Network Discovery Main Menu

11.3.1. LLDP Menu

The LLDP menu is used to configure LLDP on the switch, globally and per port, to exchange LLDP information with neighbors, and to view LLDP information and statistics.

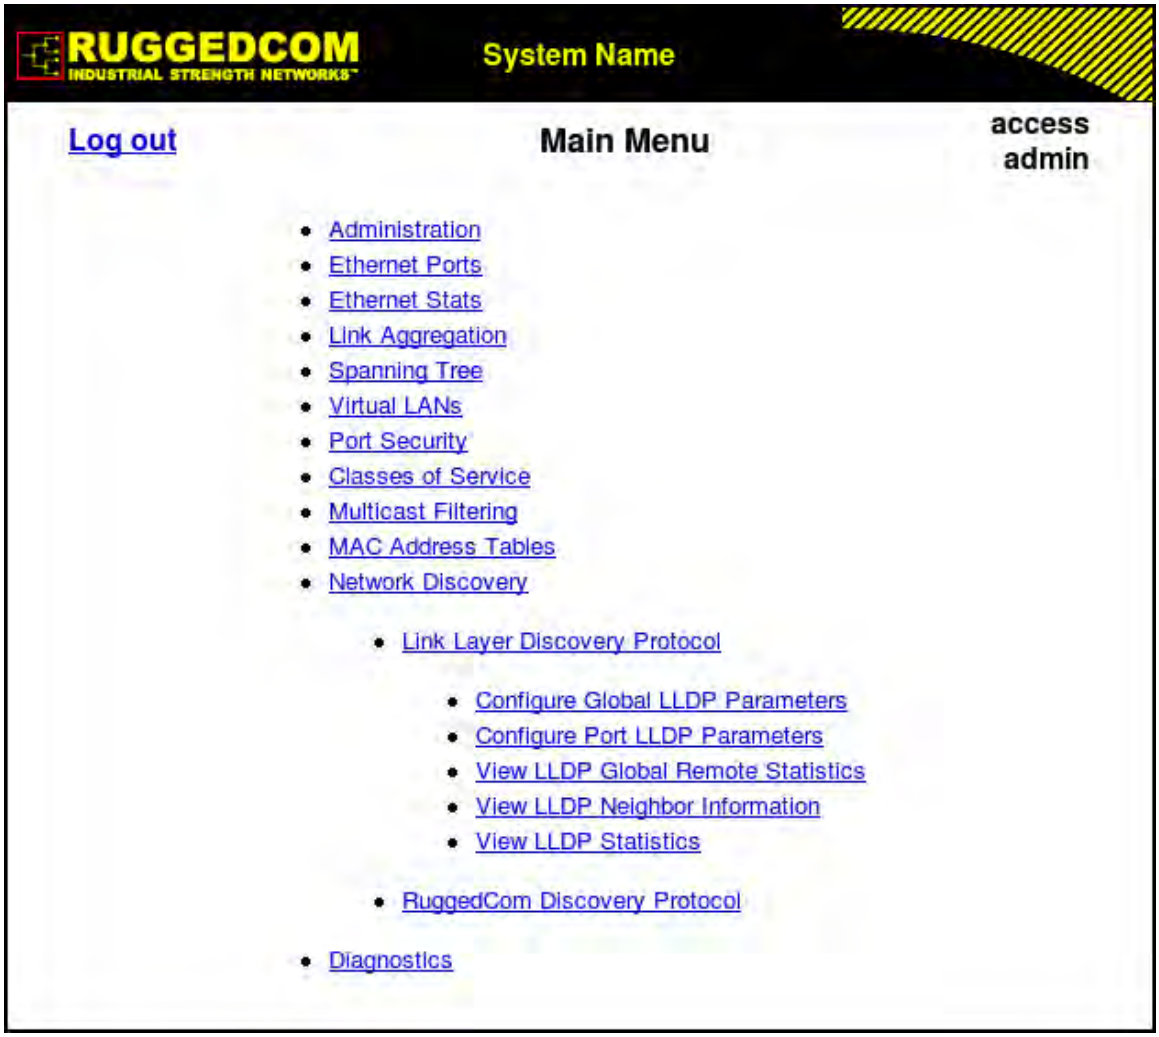

Figure 11.2. Network Discovery Menu

11.3.1.1. Global LLDP Parameters

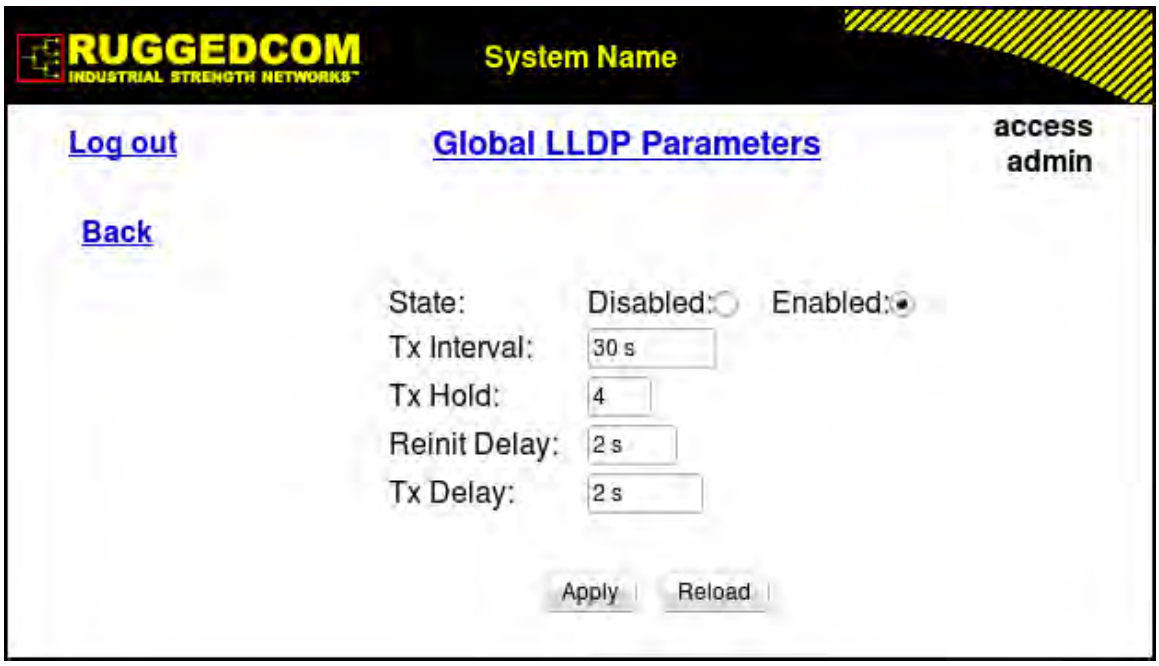

Figure 11.3. Global LLDP Parameters Form

State

```
Synopsis: { Disabled, Enabled }
Default: Enabled
```

Enables the LLDP protocol. Note that LLDP is enabled on a port when LLDP is enabled globally and along with enabling per port setting in Port LLDP Parameters menu.

Tx Interval

```
Synopsis: 5 to 32768
Default: 30 s
```

The interval at which LLDP frames are transmitted on behalf of this LLDP agent.

Tx Hold

```
Synopsis: 2 to 10
Default: 4
```

The multiplier of the Tx Interval parameter that determines the actual time-to-live (TTL) value used in a LLDPDU. The actual TTL value can be expressed by the following formula:

TTL = MIN(65535, (Tx Interval * Tx Hold))

Reinit Delay

```
Synopsis: 1 to 10
Default: 2 s
```

The delay in seconds from when the value of Admin Status parameter of a particular port becomes 'Disabled' until re-initialization will be attempted.

Tx Delay

```
Synopsis: 1 to 8192
Default: 2 s
```

The delay in seconds between successive LLDP frame transmissions initiated by value or status

changed. The recommended value is set according to the following formula:

1 <= txDelay <= (0.25 * Tx Interval)

11.3.1.2. Port LLDP Parameters

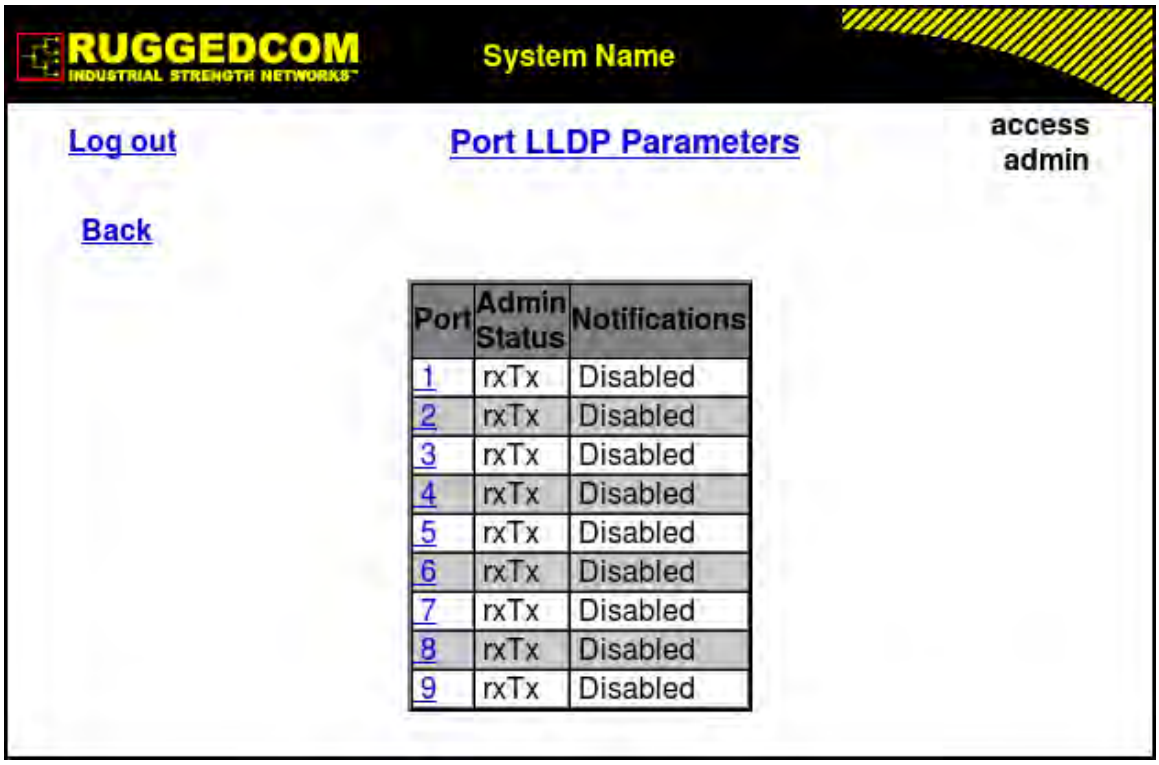

Figure 11.4. Port LLDP Parameters Table

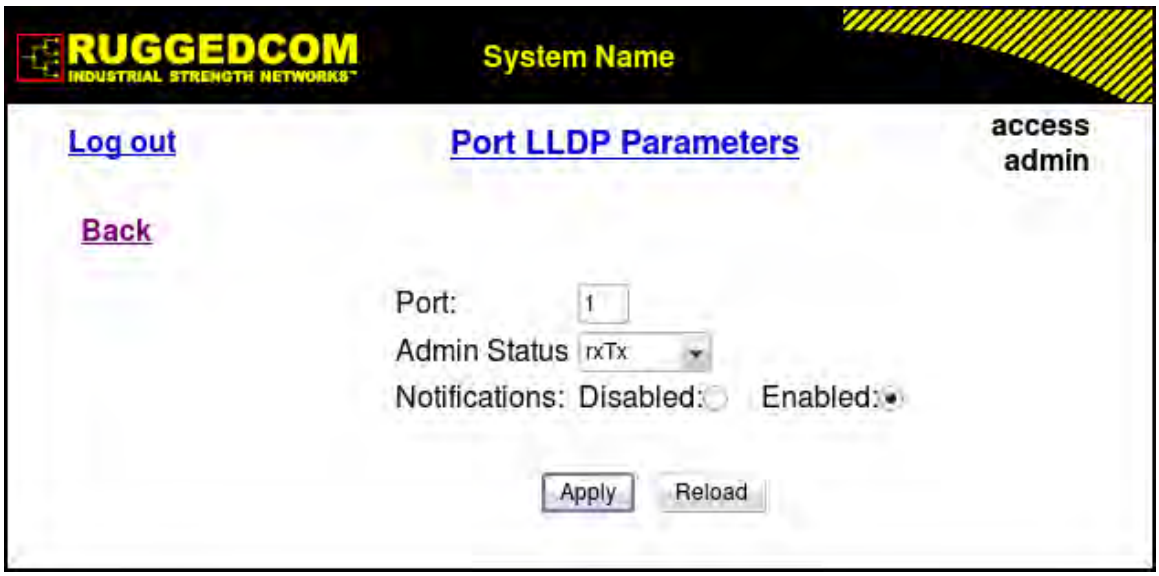

Figure 11.5. Port LLDP Parameters Form

Port

```
Synopsis: 1 to 9
Default: 1
```

The port number as seen on the front plate silkscreen of the switch.

Admin Status

```
Synopsis: { rxTx, txOnly, rxOnly, Disabled }
Default: rxTx
```

- rxTx: the local LLDP agent can both transmit and receive LLDP frames through the port.
- txOnly: the local LLDP agent can only transmit LLDP frames.
- rxOnly: the local LLDP agent can only receive LLDP frames.
- disabled: the local LLDP agent can neither transmit nor receive LLDP frames.

Notifications

```
Synopsis: { Disabled, Enabled }
Default: Disabled
```

Enabling notifications will allow the LLDP agent to send notifications and generate alarms for the port.

11.3.1.3. LLDP Global Remote Statistics

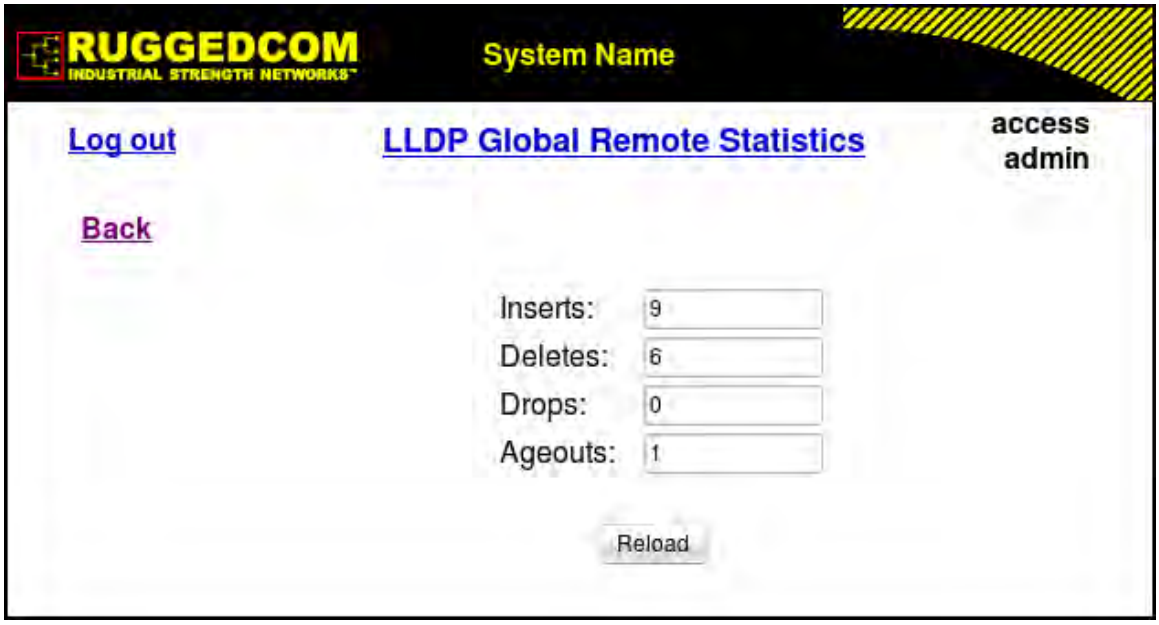

Figure 11.6. LLDP Global Remote Statistics Form

Inserts

```
Synopsis: 0 to 4294967295
```

The number of times an entry was inserted into the LLDP Neighbor Information Table.

Deletes

```
Synopsis: 0 to 4294967295
```

The number of times an entry was deleted from the LLDP Neighbor Information Table.

Drops

Synopsis: 0 to 4294967295

The number of times an entry was deleted from the LLDP Neighbor Information Table because the information timeliness interval has expired.

Ageouts

Synopsis: 0 to 4294967295

The number of all TLVs discarded.

11.3.1.4. LLDP Neighbor Information

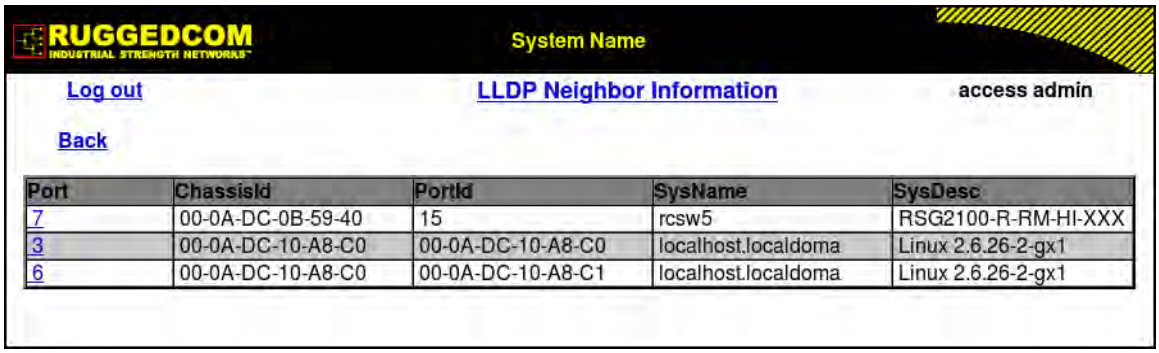

Port	ChassisId	PortId	SysName	SysDesc
7	00-0A-DC-0B-59-40	15	rcsw5	RSG2100-R-RM-HI-XXX
3	00-0A-DC-10-A8-C0	00-0A-DC-10-A8-C0	localhost.localdoma	Linux 2.6.26-2-gx1
6	00-0A-DC-10-A8-C0	00-0A-DC-10-A8-C1	localhost.localdoma	Linux 2.6.26-2-gx1

Figure 11.7. LLDP Neighbor Information Table

Port

Synopsis: 0 to 4294967295

The local port associated with this entry.

ChassisId

Synopsis: Any 19 characters

Chassis Id information received from a remote LLDP agent.

PortId

Synopsis: Any 19 characters

Port Id information received from a remote LLDP agent.

SysName

Synopsis: Any 19 characters

System Name information received from a remote LLDP agent.

SysDesc

Synopsis: Any 19 characters

System Descriptor information received from a remote LLDP agent.

11.3.1.5. LLDP Statistics

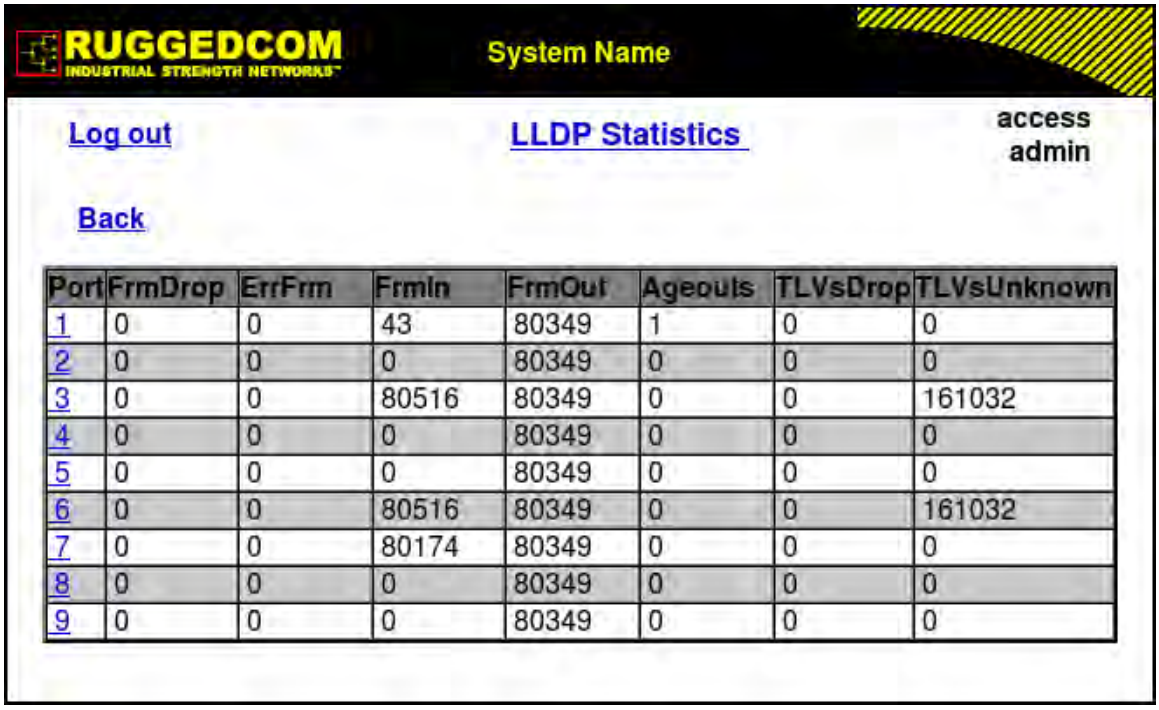

Port	FrmDrop	ErrFrm	FrmIn	FrmOut	Ageouts	TLVsDrop	TLVsUnknown
1	0	0	43	80349	1	0	0
2	0	0	0	80349	0	0	0
3	0	0	80516	80349	0	0	161032
4	0	0	0	80349	0	0	0
5	0	0	0	80349	0	0	0
6	0	0	80516	80349	0	0	161032
7	0	0	80174	80349	0	0	0
8	0	0	0	80349	0	0	0
9	0	0	0	80349	0	0	0

Figure 11.8. LLDP Statistics Table

Port

`Synopsis: 1 to 9`

The port number as seen on the front plate silkscreen of the switch.

FrmDrop

`Synopsis: 0 to 4294967295`

The number of all LLDP frames discarded.

ErrFrm

`Synopsis: 0 to 4294967295`

The number of all LLDPDUs received with detectable errors.

FrmIn

`Synopsis: 0 to 4294967295`

The number of all LLDPDUs received.

FrmOut

`Synopsis: 0 to 4294967295`

The number of all LLDPDUs transmitted.

Ageouts

`Synopsis: 0 to 4294967295`

The number of times that a neighbor's information has been deleted from the LLDP remote system MIB because the txinfoTTL timer has expired.

TLVsDrop

`Synopsis: 0 to 4294967295`

The number of all TLVs discarded.

TLVsUnknown

Synopsis: 0 to 4294967295

The number of all TLVs received on the port that are not recognized by the LLDP local agent.

11.3.2. RCDP Configuration

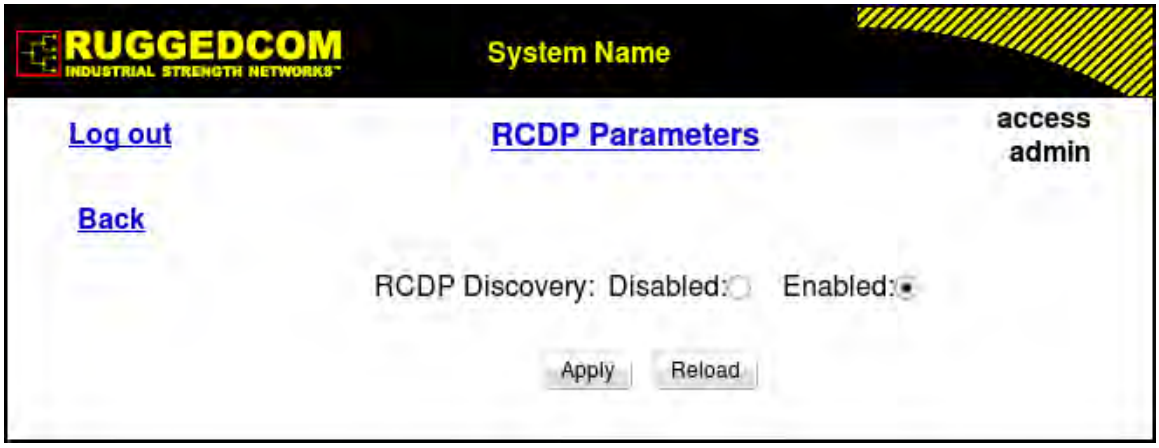

Figure 11.9. LLDP Neighbor Information Table

RCDP Discovery

Synopsis: { Disabled, Enabled }

Default: Enabled

Disables/Enables Device Discovery through RuggedCom Proprietary RCDP.

12. Diagnostics

ROS® provides the following diagnostics features:

- Alarm System to view and clear alarms.
- Viewing and clearing the system log.
- Viewing CPU diagnostics.
- Viewing the product information.
- Loading the factory default configuration.
- Resetting the device.

The Diagnostics menu is accessible from the main menu:

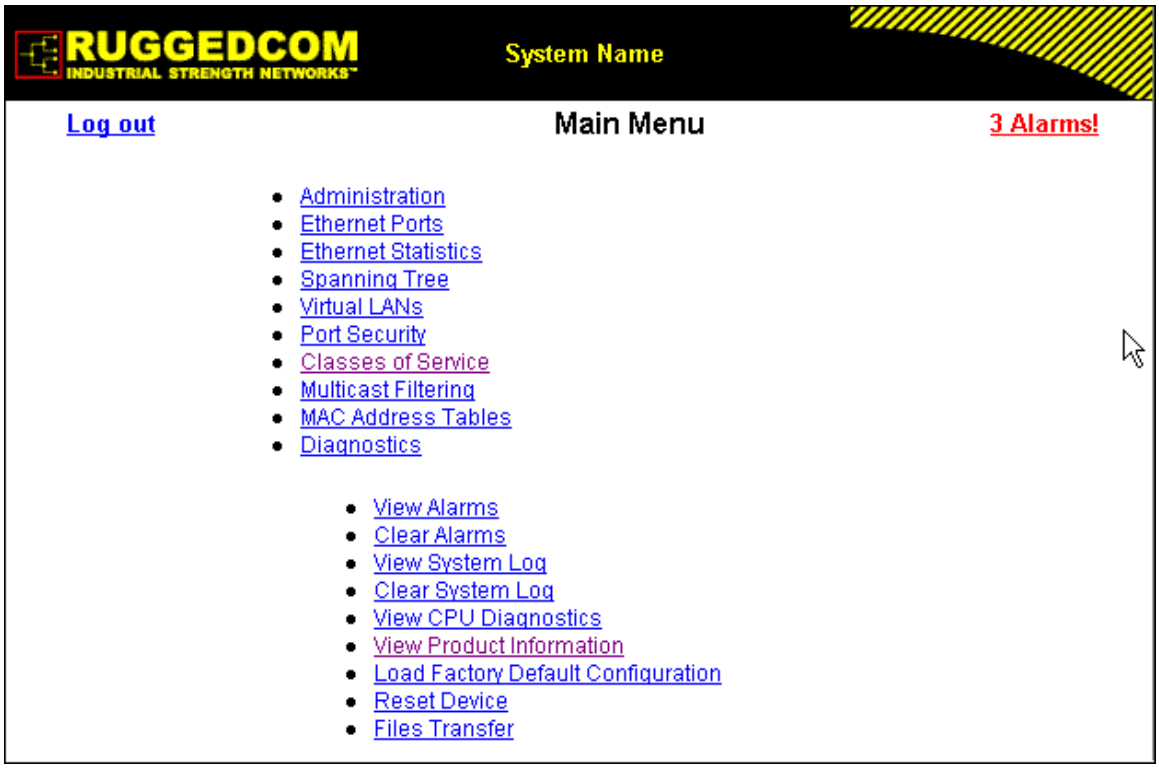

Figure 12.1. Diagnostics Menu

12.1. Using the Alarm System

Alarms are the occurrence of events of interest that are logged by the device. If alarms have occurred, the device will indicate the number of alarms in the top right corner of all menu screens.

There are two broad types of alarms - active and passive alarms.

12.1.1. Active Alarms

Active alarms are ongoing. They signify states of operation that are not in accordance with normal operation. Examples of active alarms include links that should be up but are not or error rates that are continuously exceeding a certain threshold.

Active alarms are removed (cleared) either by solving the original cause of the alarm or by explicitly clearing the alarm itself.

12.1.2. Passive Alarms

Passive alarms are historic in nature. They signify events that represented abnormal conditions in the past, and do not affect the current operational status. Examples of passive alarms include authentication failures or error rates that temporarily exceeded a certain threshold.

Passive alarms are cleared through the Clear Alarms option under the diagnostics menu. RMON generated alarms are passive.

12.1.3. Alarms and the Critical Failure Relay

All active alarms will immediately de-energize the critical fail relay (thus signifying a problem). The relay will be re-energized when the last outstanding active alarm is cleared.

Note

Alarms are volatile in nature. All alarms (active and passive) are cleared at startup.

12.1.4. Configuring Alarms

ROS provides a means for selectively configuring alarms in fine-grained detail. Some notes on alarm configuration in ROS:

- Alarms at levels CRITICAL or ALERT are not configurable nor can they be disabled.
- The "Level" field is read-only; the preconfigured alarm level is not a configurable option.
- Alarms cannot be added to or deleted from the system.
- Alarm configuration settings changed by a user will be saved in the configuration file.
- The "alarms" CLI command lists all alarms - configurable and non-configurable.

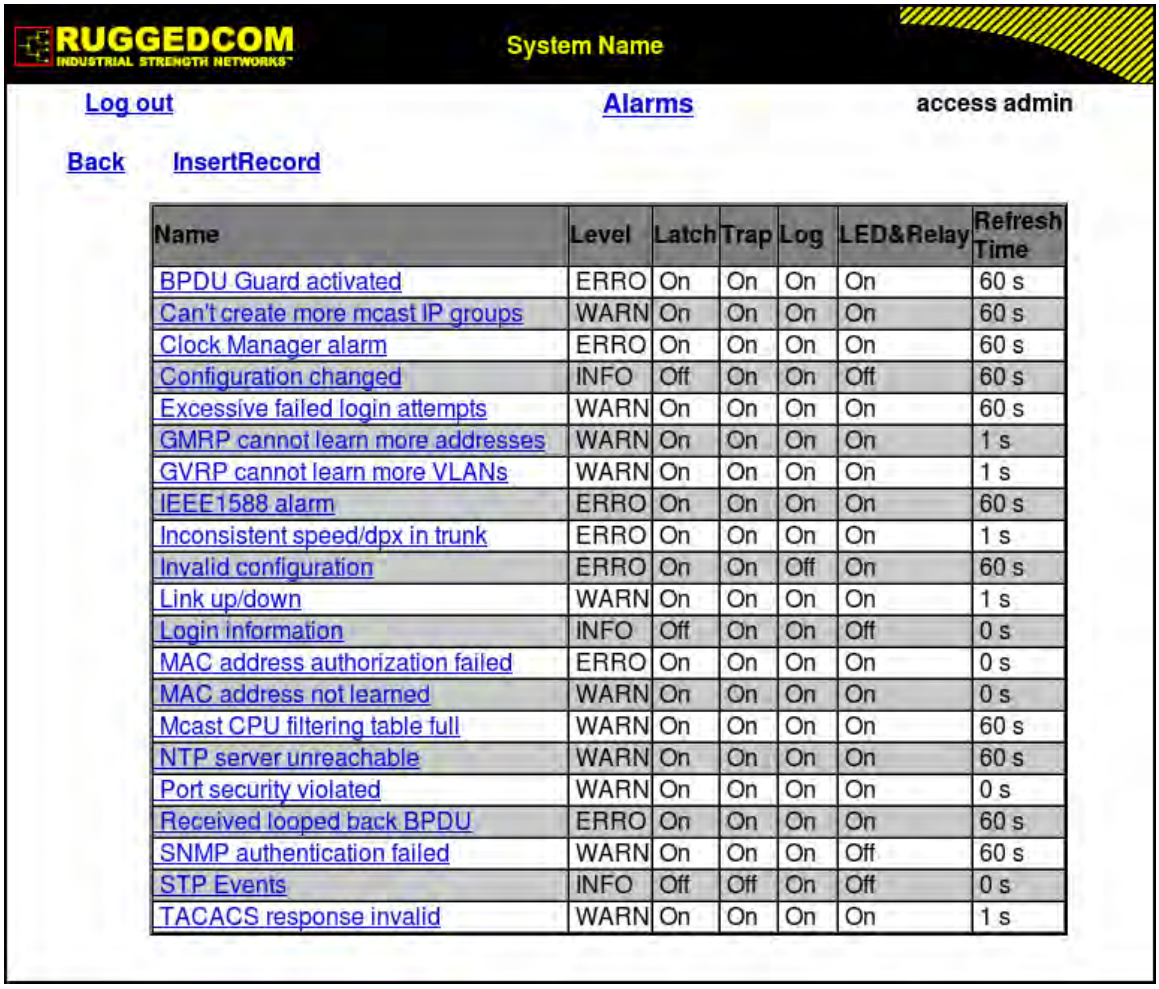

Name	Level	Latch	Trap	Log	LED&Relay	Refresh Time
BPDU Guard activated	ERRO	On	On	On	On	60 s
Can't create more mcast IP groups	WARN	On	On	On	On	60 s
Clock Manager alarm	ERRO	On	On	On	On	60 s
Configuration changed	INFO	Off	On	On	Off	60 s
Excessive failed login attempts	WARN	On	On	On	On	60 s
GMRP cannot learn more addresses	WARN	On	On	On	On	1 s
GVRP cannot learn more VLANs	WARN	On	On	On	On	1 s
IEEE1588 alarm	ERRO	On	On	On	On	60 s
Inconsistent speed/dpx in trunk	ERRO	On	On	On	On	1 s
Invalid configuration	ERRO	On	On	Off	On	60 s
Link up/down	WARN	On	On	On	On	1 s
Login information	INFO	Off	On	On	Off	0 s
MAC address authorization failed	ERRO	On	On	On	On	0 s
MAC address not learned	WARN	On	On	On	On	0 s
Mcast CPU filtering table full	WARN	On	On	On	On	60 s
NTP server unreachable	WARN	On	On	On	On	60 s
Port security violated	WARN	On	On	On	On	0 s
Received looped back BPDU	ERRO	On	On	On	On	60 s
SNMP authentication failed	WARN	On	On	On	Off	60 s
STP Events	INFO	Off	Off	On	Off	0 s
TACACS response invalid	WARN	On	On	On	On	1 s

Figure 12.2. Alarm Configuration Table

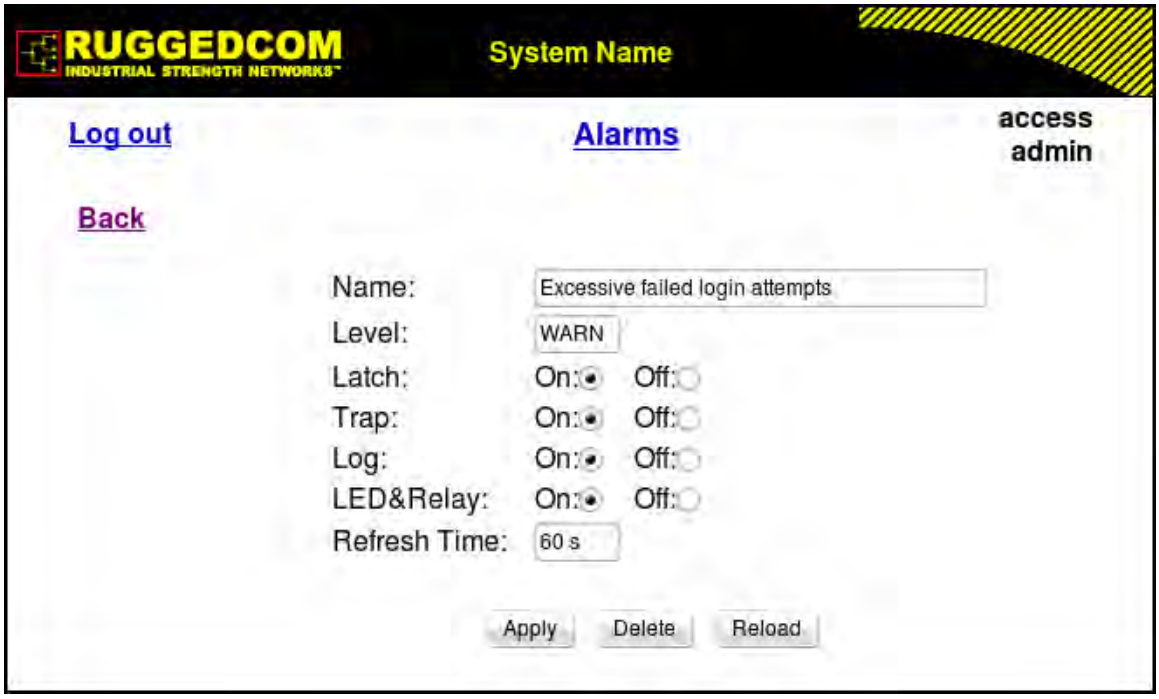

Figure 12.3. Alarm Configuration Form

Name

```
Synopsis: Any 34 characters
Default: sys_alarm
```

The alarm name (e.g. as obtained via CLI:"alarms")

Level

```
Synopsis: { EMRG, ALRT, CRIT, ERRO, WARN, NOTE, INFO, DEBG }
```

Severity level of the alarm:

- EMERG - The device has had a serious failure that caused a system reboot.
- ALERT - The device has had a serious failure that did not cause a system reboot.
- CRITICAL - The device has a serious unrecoverable problem.
- ERROR - The device has a recoverable problem that does not seriously affect operation.
- WARNING - Possibly serious problem affecting overall system operation.
- NOTIFY - Condition detected that is not expected or not allowed.
- INFO - Event which is a part of normal operation, e.g. cold start, user login etc.
- DEBUG - Intended for factory troubleshooting only.

Latch

```
Synopsis: { On, Off }
Default: Off
```

Enables latching occurrence of this alarm in the Alarms Table.

Trap

```
Synopsis: { On, Off }
Default: Off
```

Enables sending an SNMP trap for this alarm.

Log

```
Synopsis: { On, Off }
Default: Off
```

Enables logging the occurrence of this alarm in syslog.txt.

LED & Relay

```
Synopsis: { On, Off }
Default: Off
```

Enables LED and fail-safe relay control for this alarm. If latching is not enabled, this field will remain disabled.

Refresh Time

```
Synopsis: 0 s to 60 s
Default: 60 s
```

Refreshing time for this alarm.

12.1.5. Viewing and Clearing Alarms

Alarms are displayed in the order in which they occurred, even if the real time clock was incorrect at the time of the alarm.

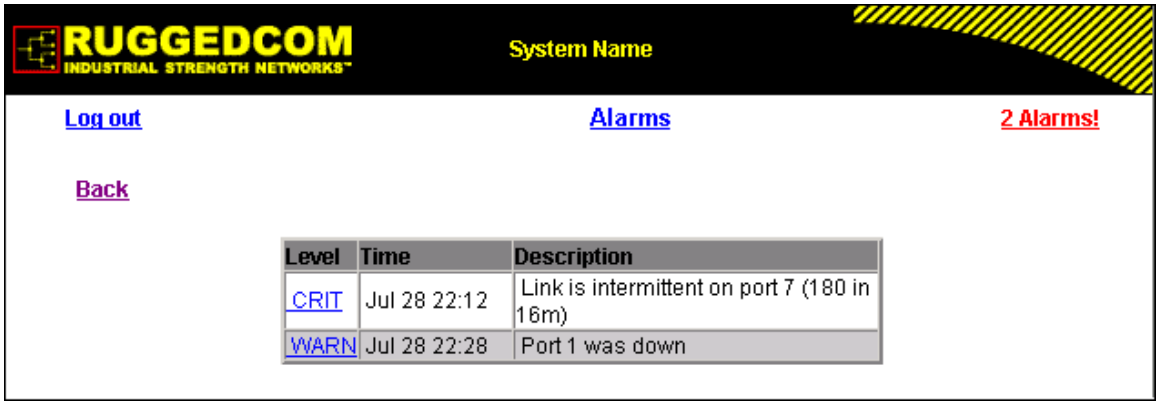

Figure 12.4. Alarm Table

Level

```
Synopsis: { EMRG, ALRT, CRIT, ERRO, WARN, NOTE, INFO, DEBG }
```

Severity level of alarm - refer to Level, above, for a detailed breakdown of the levels.

Time

```
Synopsis: MMM DD HH:MM
```

Time of first occurrence of the alarm.

Description

```
Synopsis: Any 127 characters
```

Description of the alarm; gives details about the frequency of the alarm if it has occurred again since the last clear.

Alarms can be cleared from the Clear Alarms option.

12.2. Viewing CPU Diagnostics

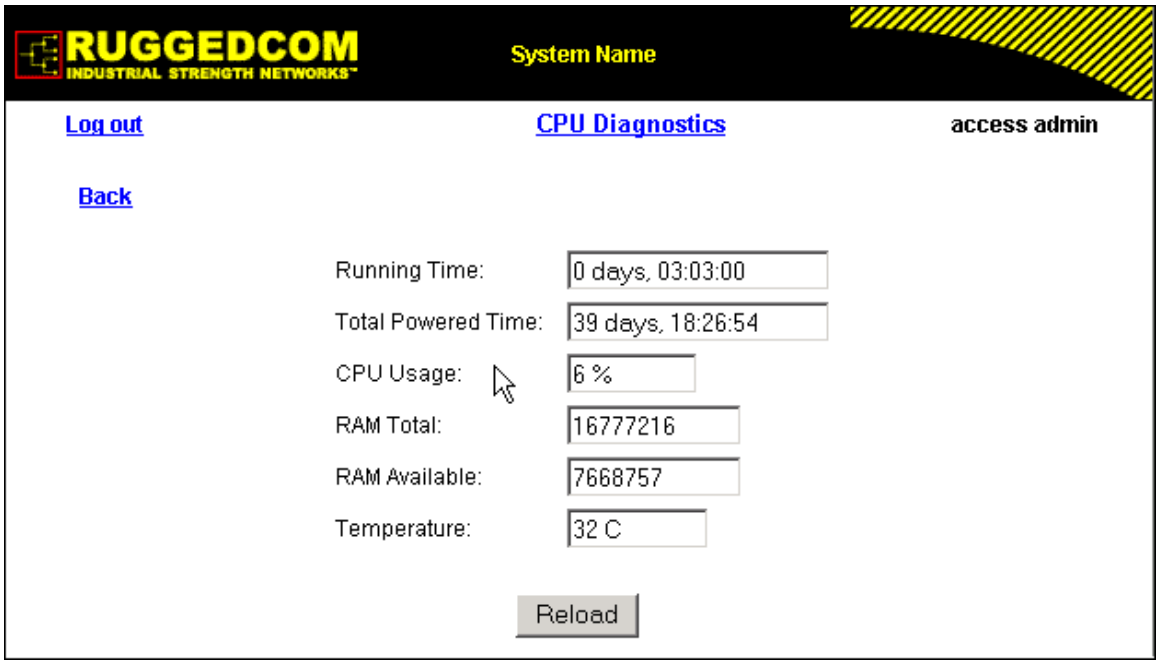

Figure 12.5. CPU Diagnostics Form

Running Time

`Synopsis: DDDD days, HH:MM:SS`

The length of time since the device was last powered on.

Total Powered Time

`Synopsis: DDDD days, HH:MM:SS`

The cumulative powered up time of the device.

CPU Usage

`Synopsis: 0 to 100`

The percentage of available CPU cycles used for device operation as measured over the last second.

RAM Total

`Synopsis: 0 to 4294967295`

The total number of bytes of RAM in the system.

RAM Available

`Synopsis: 0 to 4294967295`

The total number of bytes of RAM still available.

Temperature

`Synopsis: -32768 to 32767 C`

The temperature of the CPU board.

12.3. Viewing and Clearing the System Log

The system log records various events including reboots, user sign-ins, alarms and configuration saves.

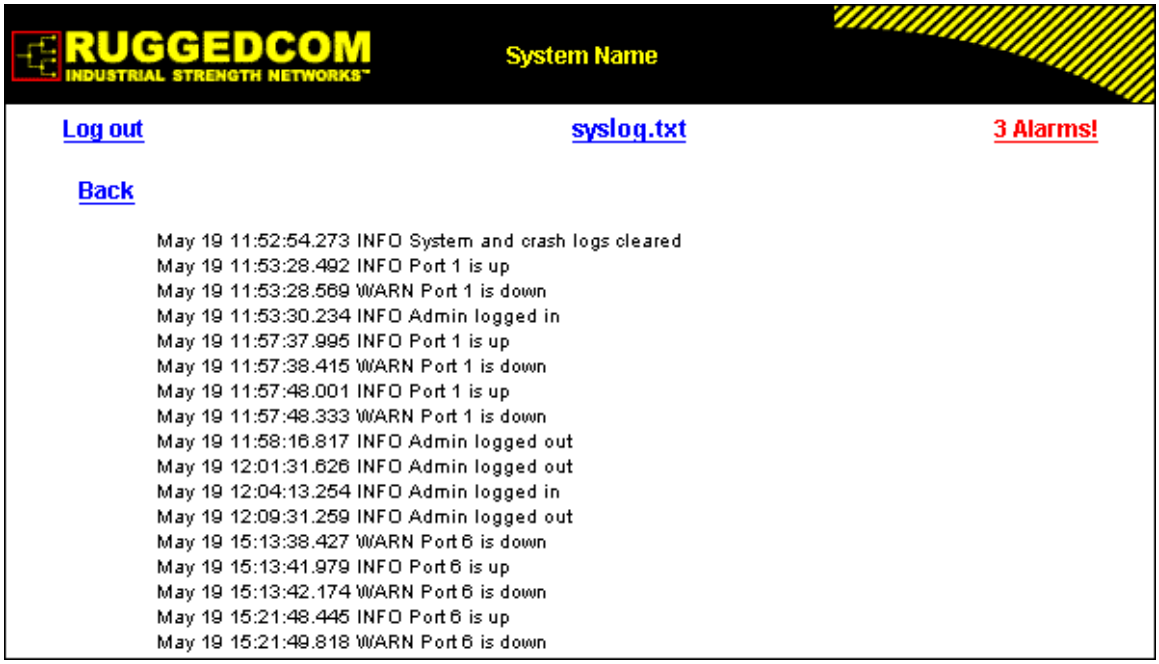
RUGGEDCOM
INDUSTRIAL STRENGTH NETWORKS
System Name
Log out
syslog.txt
3 Alarms!
Back
May 19 11:52:54.273 INFO System and crash logs cleared
May 19 11:53:28.492 INFO Port 1 is up
May 19 11:53:28.569 WARN Port 1 is down
May 19 11:53:30.234 INFO Admin logged in
May 19 11:57:37.995 INFO Port 1 is up
May 19 11:57:38.415 WARN Port 1 is down
May 19 11:57:48.001 INFO Port 1 is up
May 19 11:57:48.333 WARN Port 1 is down
May 19 11:58:16.817 INFO Admin logged out
May 19 12:01:31.626 INFO Admin logged out
May 19 12:04:13.254 INFO Admin logged in
May 19 12:09:31.259 INFO Admin logged out
May 19 15:13:38.427 WARN Port 6 is down
May 19 15:13:41.979 INFO Port 6 is up
May 19 15:13:42.174 WARN Port 6 is down
May 19 15:21:48.445 INFO Port 6 is up
May 19 15:21:49.818 WARN Port 6 is down

Figure 12.6. Viewing the System Log

The system log will continue to accumulate information until it becomes full. There is enough room in the file to accumulate logs for months or years under normal operation.

The Clear System Log option will clear the system log. Clearing the log is recommended after a firmware upgrade.

12.4. Viewing Product Information

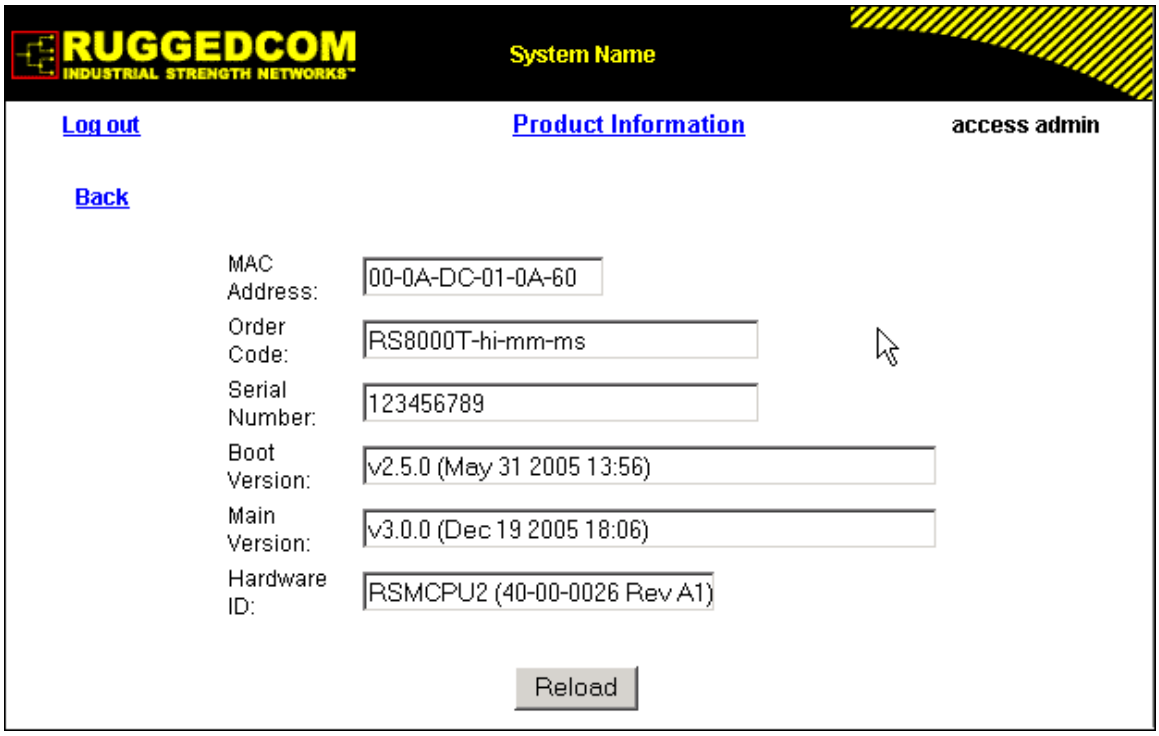

Figure 12.7. Product Information Form

MAC Address

`Synopsis: ##-##-##-##-##-## where ## ranges 0 to FF`

Shows the unique MAC address of the device.

Order Code

`Synopsis: 31 characters`

Shows the order code of the device.

Serial Number

`Synopsis: 31 characters`

Shows the serial number of the device.

Boot Version

`Synopsis: 47 characters`

Shows the version and the build date of the boot loader software.

Main Version

`Synopsis: 47 characters`

Shows the version and build date of the main operating system software.

Hardware ID

`Synopsis: { 47 characters }`

Shows the type, part number, and revision level of the hardware.

12.5. Loading Factory Default Configuration

The Load Factory Defaults menu is used to reset the unit's configuration to its factory default. Optionally, it is possible to exclude parameters that affect basic connectivity and SNMP management from the reset in order to be able to remain in communication with the device. Specifically, configuration items in the following categories are **not** affected by a selective configuration reset:

- IP Interfaces
- IP Gateways
- SNMP Users
- SNMP Security to Group Maps
- SNMP Access

The menu presents a choice of whether to reset all or only the selected set of configuration parameters to their factory default values:

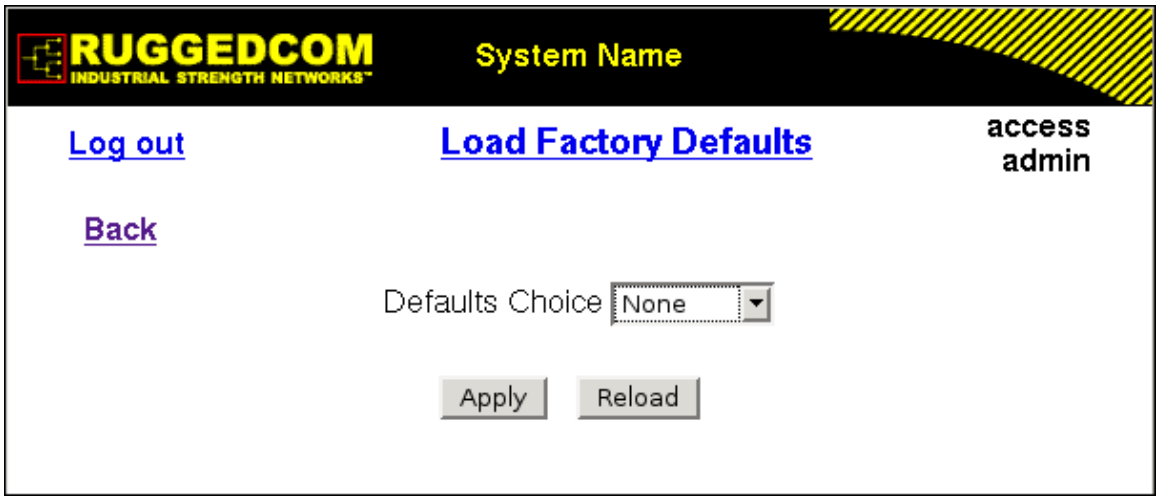

Figure 12.8. Load Factory Defaults Dialog

Defaults Choice

`Synopsis: { None, Selected, All }`

This parameter allows the user to choose to load defaults to Selected tables (i.e. excluding those listed above), which would preserve configuration of the tables that are critical for basic communication and switch management applications, or to force All tables to default settings.

Note

It is possible to explicitly reset configuration items in the exceptional categories listed above to their default values by using the "sql" command. Please refer to the section entitled: "Upgrading Firmware and Managing Configurations".

12.6. Resetting the Device

This operation will close all open Telnet connections and warm-start the device after the user has confirmed the reset operation from the Reset Device option.

Figure 12.9. Reset Device Dialog

13. Using the CLI Shell

ROS® Command Line Interface (CLI) support enables:

- Execution of commands from a CLI shell.
- Remote execution of commands using RSH or SSH.
- Switching between the CLI shell and the menu system.

Note

Different commands may be available to users at different login session security levels (guest, operator or administrator).

The ROS CLI shell may be accessed from a terminal session to the device. A terminal session may be established in one of three ways:

- Direct cable, via RS-232.
- Remote via RSH.
- Remote via SSH.

When a terminal session is first established to the ROS device, the user interface presented will be the full-screen menu interface. Please refer to Section Section 1.1, “The ROS® User Interface” for more detail on the menu interface.

The Command Line Interface (CLI) shell may be accessed from any menu by pressing <Ctrl-S>. Any menu operation in progress, such as changing a configuration parameter, will be terminated. You may return to the menu system by pressing <Ctrl-S> again or by entering “exit<CR>” at the shell prompt.

This chapter describes a selection of the most useful commands in detail. For a complete list of available commands, please refer to Appendix E, *Command Line Listing*.

13.1. Summary Of CLI Commands available in ROS®

Type “help<CR>” to see the list of commands available at the current session access level.

```
>help

alarms          Displays list of available alarms.
clearalarms     Clears all alarms.
clearethstats   Clears statistics for specified Ethernet port(s).
clearlogs       Clears the system and crash logs.
cls             Clears the screen
delay           Pause a specified number of milliseconds.
dir             Prints file directory listing.
echo            Echoes the specified message to the screen.
exit            Terminate this command line session
help            Print listing of all commands
ipconfig        Displays IP configuration.
login           Login to the shell i.e. set the access level
logout          Logout of the shell
ping            Pings specified IP address
reset           Perform a 'hard' reset of the switch
resetport       Resets specified switch port(s).
rmon            Displays names of RMON alarm eligible objects.
route           Displays gateway configuration.
spuriouscount   Displays User Spurious Interrupt count.
sql             'SQL' like commands for setting/viewing system parameters
sudp            Sends UDP msgs to 192.168.0.255
telnet          Telnet to the server with specified IP address
tftp            TFTP client executes command on server specified by IP address.
trace           trace command
type            Displays the contents of a text file.
version         Prints software versions.
xmodem          Upload or download a file to the switch.
```

Figure 13.1. Displaying The List Of Available Commands

13.2. Obtaining Help For A Command

Help related to the usage of a particular command may be obtained by entering “help command name <CR>” at the shell prompt.

```
>help type
Displays the contents of a text file.
Enter 'dir' for a directory listing of files.

TYPE filename
```

Figure 13.2. Displaying Help For A Command

13.3. Viewing Files

RuggedCom devices maintain a number of volatile and non-volatile files. These files can aid in the resolution of problems and serve as a useful gauge of the device’s health.

13.3.1. Listing Files

Enter “dir<CR>” to obtain a complete list of files and a description of each.

Note

Each file has associated attributes, as described under the Attr column in “dir” command. Files marked “R” are readable, i.e. may be uploaded by the user. Files marked “W” are writable,

i.e. may be modified (downloaded) by the user. Files marked "B" are binary files, i.e. may be upgraded by the user.

The most useful files include config.csv, crashlog.txt and syslog.txt. These files may be viewed by using the "type" command, specifying the desired filename.

```
>dir
Directory of RuggedSwitch
--------------------------------------------------------------------------------
Free files:   21 of 32
Free handles: 31 of 32
Free blocks:  1024 of 1024
Block size:   4096
--------------------------------------------------------------------------------
Filename             Size Hdls Blks Attr Description
--------------------------------------------------------------------------------
dir.txt                 0    1    1 R   Listing of files and attributes.
boot.bin           342930    0    0 RWB Boot firmware
main.bin          1424310    0    0 RWB Operating system firmware
fpga.xsvf           55921    0    0 RWB FPGA programming file binary file
factory.txt           161    0    0 RW  Factory data parameters
config.csv           8555    0    0 RW  System settings
config.bak           8555    0    0 RW  System settings backup
crashlog.txt            0    0    0 RW  Log of debilitating system events
syslog.txt           3105    0    0 RW  Log of system events
sdram.bin        16777216    0    0 R B Image of entire SDRAM memory
flash.bin         4194304    0    0 R B Image of entire Flash memory
--------------------------------------------------------------------------------
```

Figure 13.3. Displaying The Directory Of A ROS Device

13.3.2. Viewing and Clearing Log Files

The crashlog.txt and syslog.txt files contain historical information about events that have occurred.

The crashlog.txt file will contain debugging information related to problems that might have resulted in unplanned restarts of the device or which may effect the device operation. A file size of 0 bytes indicates that no untoward events have occurred.

The syslog.txt file contains a record of significant events including startups, configuration modifications, firmware upgrades and database re-initializations due to feature additions. Syslog.txt file will accumulate information until it fills, holding approximately 3 megabytes of characters.

The "clearlogs" command resets these logs. It is recommended to run "clearlogs" command after every firmware upgrade.

13.4. Pinging a Remote Device

The "ping" command sends an ICMP echo request to a remotely connected device. For each reply received, the round trip time is displayed.

The command, "ping <IP address>", will send a small number of pings to the device with this IP address and display the results. The ping command can be used to verify connectivity to the next connected device. It is a useful tool for testing commissioned links. This command also includes the ability to send a specific number of pings with a specified time for which to wait for a response.

The specification of a large number of pings and a short response time can “flood” a link, stressing it more than a usual ping sequence. The command “ping 192.168.0.1 500 2” can be used to issue 500 pings, each separated by two milliseconds to the next device. If the link used is of high quality, then no pings should be lost and the average round trip time should be small.

Note

The device to be pinged must support ICMP echo. Upon commencing the ping, an ARP request for the MAC address of the device is issued. If the device to be pinged is not on the same network as the device pinging the other device, the default gateway must be programmed.

13.5. Tracing Events

The CLI trace command provides a means to trace the operation of various protocols supported by the device. Trace provides detailed information including RSTP packet decodes, IGMP activity and MAC address displays.

Note

Tracing has been designed to provide detailed information to expert users. Note that all tracing is disabled upon device startup.

In order to display the current trace settings and discover the systems that are being traced, enter the CLI command “trace ?”.

```
trace ?
Supported commands:
noclear          Starts the log without clearing it first
alloff           Disables all trace subsystems from tracing
allon            Enables all flags in all trace subsystems
stp              Traces STP operations
link             Displays switch fabric statistics
mac              Displays MAC Events
forward          Forwards trace messages to an IP:UDP address
igmp             Displays IGMP Snooping events
gvrp             Displays GVRP events
webs             Traces Web Server connections
dhcpra           Traces DHCP Relay Agent
802.1X           Traces 802.1X PAE
ip               Traces IP communications

Enter "trace command ?" for more information on a particular command.

STP  : Logging all conditions on port(s) 1-10
LINK : Logging is disabled
MAC  : Logging is disabled
FORW : IP: 0.0.0.0  UDP: 0 (OFF)
IGMP : Logging is disabled
GVRP : Logging is disabled
WEBS : Logging is disabled
DHCPRA : Logging is disabled
802.1X : Logging is disabled
IP   : Logging is disabled
```

Figure 13.4. Displaying Trace Settings

13.5.1. Enabling Trace

Tracing can be enabled on a per subsystem basis. Obtain detailed information about individual subsystems by entering “trace subsystem_name ?<CR>”. Some subsystems offer a mechanism to enable tracing only on certain ports.

```
>trace stp ?
trace stp syntax:
  stp [-|+] [all] [verbose] [packets] [timers] [actions]
            [decodes] [ports[port_number|all]]
STP  : Logging is disabled

>trace stp all
STP  : Logging all conditions on port(s) 1-16

>trace link ?
trace link syntax
  link changes | stats | allon | alloff | statsonce
LINK : Logging is disabled

>trace link changes
LINK : changes
>
```

Figure 13.5. Enabling Trace

13.5.2. Starting Trace

To start trace, enter “trace<CR>”. All historical trace messages may be displayed using “trace noclear<CR>”. Since this may include many messages, it may be more desirable to use the “trace clear<CR>” command instead. This command will automatically clear the trace buffer as it starts the trace.

```
>trace stp - all

STP  : Logging is disabled
>trace stp decodes

STP  : Logging decodes
>trace stp port 7

STP  : Logging decodes on port(s) 7

> trace link changes
LINK : changes

>trace

Log has been cleared
009.445 IGMP TX General Query, VLAN    1, gr. 000.000.000.000,
             to ports ALL VLAN PORTS
010.543 LINK Link 7 has risen.
000.550 RSTP TX port 7 RST BPDU: TCack 0 agg 1 lrn 0 fwd 0 role DP prop 1 TC 0
            root 32768/0adc001000 cst 38, brdg 32768/0adc005000, prt 128/7
            age 2.00, maxage 20, hello 2, fwddelay 15 V1Length 0
000.557 RSTP RX port 7 RST BPDU: TCack 0 agg 1 lrn 0 fwd 0 role DP prop 1 TC 0
            root 32768/0adc004000 cst 0, brdg 32768/0adc004000, prt 128/14
            age 0.00, maxage 20, hello 2, fwddelay 15 V1Length 0
```

Figure 13.6. Starting Trace

Note

The trace package includes the “forward” subsystem, a remote reporting facility intended to be used only under the direction of RuggedCom service personnel.

13.6. Viewing DHCP Learned Information

The CLI command “ipconfig<CR>” will provide the current IP address, subnet mask and default gateway. This command provides the only way of determining these values when DHCP is used.

13.7. Executing Commands Remotely Through RSH

The Remote Shell (RSH) facility can be used from a workstation to cause the product to act upon commands as if they were entered at the CLI prompt. The syntax of the RSH command is usually of the form:

```
rsh ipaddr –l auth_token command_string
```

where:

- `ipaddr` = The address or resolved name of the RuggedCom device.
- `auth_token` = The authentication token, which for ROS rsh is the user name (guest, operator, or admin) and corresponding password separated by a comma. For example, to run a command as user - "admin" with password - "secret", the token would be: "admin,secret".
- `command_string` = The ROS shell command to execute.

The access level (corresponding to the user name) selected must support the given command.

Any output from the command will be returned to the workstation submitting the command. Commands that start interactive dialogs (such as trace) cannot be used.

13.8. Resetting the Device

The CLI command “reset<CR>” can be used to reset the device.

14. Firmware Upgrade and Configuration Management

ROS® provides flexible, powerful mechanisms for the bulk update and backup of system firmware and of the configuration database. The ROS firmware and configuration database are represented as files in the internal file system, and bulk update and backup consist of simply transferring files to and from the ROS device, by one of the several means provided.

ROS also implements an SQL command language in order to provide the flexibility and power of a database model when configuring ROS-based devices.

Files Of Interest

The files in ROS that may be updated and backed up are described below:

- **main.bin**: the main ROS application firmware image – Upgrades to ROS are made via updates to this file.
- **boot.bin**: the boot loader firmware image – In normal practice, the boot loader does not require updating.
- **fpga.xsvf**: the FPGA firmware binary image – also not normally updated.
- **config.csv**: the complete configuration database, in the form of a comma-delimited ASCII text file.

File Transfer Mechanisms

Several mechanisms are available to transfer these files to and from a ROS-based device:

- **Xmodem** using the ROS CLI over a (telnet or RS232) console session.
- **TFTP client** (using the ROS CLI in a console session and a remote TFTP server).
- **TFTP server** (from a remote TFTP client).
- **SFTP** (secure FTP over SSH, from a remote SFTP client).

Console Sessions

Console sessions may be established (depending on the settings in the IP Services menu) by the following means:

- **RS232** direct RS232 serial connection to the ROS device.
- **telnet** remote terminal protocol via TCP/IP (unencrypted).
- **RSH** Remote SHell, the remote login shell protocol via TCP/IP (unencrypted).
- **SSH** Secure SHell, the standard remote login shell protocol via TCP/IP – Both authentication and session are encrypted.

14.1. Upgrading Firmware

Upgrading ROS firmware may sometimes be necessary in order to take advantage of new features or bug fixes. In normal circumstances, only the main ROS application firmware is updated; the boot loader and FPGA firmware remain invariant. The main ROS application firmware image is a binary file available from RuggedCom. Please check the RuggedCom web site, www.RuggedCom.com, for the availability of updates to ROS firmware or contact RuggedCom support.

Firmware upgrades may be performed using any of the transfer methods and protocols listed under File Transfer Mechanisms, above.

Applying the Upgrade

Binary firmware images transferred to the ROS-based device are stored in non-volatile memory and require a device reset in order to take effect. The “version” ROS shell command will display any firmware updates that are pending. Currently running firmware is labeled “Current”; pending upgrades are labeled “Next”:

```
>version
Current ROS-CF52 Boot Software v2.14.0 (Sep 29 2008 13:25)
Current ROS-CF52 Main Software v3.6.0 (Oct 03 2008 09:33)
Next    ROS-CF52 Main Software v3.7.0 (Jun 02 2009 08:36)
```

ROS firmware is provided as a compressed installation image. When this compressed image is run for the first time, it decompresses itself and reinstalls the decompressed image to Flash memory. Subsequent device reboots will use the decompressed image.

Security Considerations

File transfers using methods that require ROS login authentication, namely Xmodem, SFTP, and the ROS TFTP client, are subject to the following conditions:

- transfers **from** the ROS-based device may be performed by any user with login privileges.
- transfers **to** the ROS-based device may only be performed by those with administrator privileges.

The exception is that the SFTP server does not support transmission of the firmware or configuration file using anything less than administrator privileges.

File transfers (in both directions) that make use of the ROS TFTP server do not require authentication, since TFTP does not define an authentication scheme. Instead, the TFTP server must be enabled from the IP Services Configuration Menu when it is needed.

Note

It is recommended to use the ROS TFTP server (or any TFTP server) only on a secure network, owing to TFTP’s lack of an authentication scheme. Even so, and especially in a production environment, it is also recommended to leave the TFTP server enabled for only as long as it is needed.

The following sections describe briefly how to upgrade the main application firmware using each of the mechanisms provided by ROS.

14.1.1. Upgrading Firmware Using XModem

This method requires that the binary image file of the main ROS application firmware, along with serial terminal or telnet software and the ability to do Xmodem transfers, be available on a computer with an RS232 or network connection, respectively, to the ROS device to be upgraded.

Establish a console connection with administrative privileges, either via the RS232 port or via telnet. Enter the ROS command, “xmodem receive main.bin<CR>”. When ROS responds with “Press Ctrl-X to cancel”, begin your Xmodem transmission, using the means provided by your terminal software. After the file transfer has been completed, the device will provide an indication that the file has been transferred successfully. The transcript of a sample exchange, looking at the ROS CLI, follows:

```
>xmodem receive main.bin
Press Ctrl-X to cancel
Receiving data now ...C
Received 1428480 bytes. Closing file main.bin ...
main.bin transferred successfully
```

If possible, select the “XModem 1K” protocol for transmission; otherwise, select “XModem”. The device must be reset in order for the new software to take effect. If you want to reset the device immediately, enter “reset<CR>”. The device will reboot within a few seconds.

14.1.2. Upgrading Firmware Using the ROS TFTP Server

This method requires that the binary image file of the main ROS application firmware, along with TFTP client software, be available on a computer with a network connection to the ROS device to be upgraded.

Note

*The **TFTP Server** parameter in **IP Services Configuration** controls how a TFTP client can access the device’s built-in TFTP server. A setting of “Disabled” prevents all access, “Get Only” allows retrieval of files only, and “Enabled” allows both storing and retrieval of files. Ensure that this parameter is set appropriately for the type of access you wish to perform.*

Enable TFTP transfers to the ROS device, as noted above. Begin a TFTP transfer in binary mode to the device, specifying a destination filename of “main.bin”. A TFTP client utility will provide an indication that the file was transferred properly, but it is recommended to also query the device directly in order to verify successful transfer. Establish a console session to the ROS device (using RS232, telnet, or SSH) and enter the “version” command, as described in Applying the Upgrade, above. If the transfer was successful, the version of the firmware file that was transferred will appear as the “Next” firmware version, i.e. that will appear after the next reset.

The transcript of a sample TFTP transfer, looking at a DOS/Windows CLI, follows:

```
C:\>tftp -i 10.1.0.1 put C:\files\ROD-CF52_Main_v3.7.0.bin main.bin
Transfer successful: 1428480 bytes in 4 seconds, 375617 bytes/s
```

14.1.3. Upgrading Firmware Using the ROS® TFTP Client

This method requires that the binary image file of the main ROS application firmware, along with a correctly configured TFTP server, be available on a computer with a network connection to the ROS device to be upgraded.

Identify the IP address of the host providing the TFTP server capability. Ensure that the firmware revision to be downloaded (e.g. ROS-CF52_Main_v3.7.0.bin) is present there. Establish a console

connection with administrative privileges to the ROS device to be upgraded (i.e. via RS232, telnet, or SSH). Enter the CLI shell and run the TFTP client command to receive the firmware image, for example:

```
tftp <TFTP server> get <remote filename> main.bin
```

where:

- **TFTP server** is the IP address of the TFTP server
- **remote filename** is the name of the binary image file of the main ROS application firmware residing in the TFTP server outgoing directory

Verify, as above, the successful transfer via the ROS CLI "version" command. A sample transcript from the ROS CLI:

```
>tftp 10.0.0.1 get ROS-CF52_Main_v3.7.0.bin main.bin
TFTP CMD: main.bin transfer ok. Please wait, closing file ...
TFTP CMD: main.bin loading succesful.

>version
Current ROS-CF52 Boot Software v2.14.0 (Sep 29 2008 13:25)
Current ROS-CF52 Main Software v3.6.0 (Oct 03 2008 09:33)
Next    ROS-CF52 Main Software v3.7.0 (Jun 02 2009 08:36)
```

14.1.4. Upgrading Firmware Using SFTP

This method requires that the binary image file of the main ROS application firmware, along with SFTP client software, be available on a computer with a network connection to the ROS device to be upgraded. SFTP is the Secure File Transfer Protocol (also known as the SSH File Transfer Protocol), a file transfer mechanism that uses SSH to encrypt every aspect of file transfer between a networked client and server.

Establish an SFTP connection with administrative privileges to the ROS device to be upgraded. Begin a transfer to the device, specifying a destination filename of "main.bin". An SFTP client utility will provide an indication that the file was transferred properly, but, again, it is recommended to also query the device directly in order to verify successful transfer. A sample SFTP session to upgrade the ROS main firmware image from a Linux workstation follows:

```
user@host$ sftp admin@ros_ip
Connecting to ros_ip...
admin@ros_ip's password:
sftp> put ROS-CF52_Main_v3-7-0.bin main.bin
Uploading ROS-CF52_Main_v3-7-0.bin to /main.bin
ROS-CF52_Main_v3-7-0.bin                100% 2139KB  48.6KB/s   00:44
sftp>
```

14.2. Updating Configuration

ROS maintains its complete configuration in an ASCII text file, in CSV (Comma-Separated Value) format. All configuration changes, whether they are performed using the web interface, console interface, CLI, SNMP, or SQL, are stored in this one file. The file, named **config.csv**, may be read from and written to the ROS device in all the same ways that firmware image files can, as described in the preceding sections. The configuration file may be copied from the unit and used

as a backup, to be restored at a later date. Configuration files from different units may be compared using standard text processing tools.

The transfer mechanisms supported for the update of **config.csv** are the same as for ROS firmware image files:

- **Xmodem** using the ROS CLI over a console session.
- **TFTP client** (using the ROS CLI in a console session and a remote TFTP server).
- **TFTP server** (from a remote TFTP client).
- **SFTP** (secure FTP over SSH, from a remote SFTP client).

Please refer to the preceding section, Section 14.1, “ Upgrading Firmware ”, for examples of the use of each of these mechanisms for transferring a file to a ROS device.

Configuration File Format

The format of the configuration file makes it simple to apply a wide variety of tools to the task of maintaining ROS configuration. Among the applications that may be used to manipulate ROS configuration files are:

- Any text editing program capable of reading and writing ASCII files.
- Difference/patching tools (e.g. the UN*X “diff” and “patch” command line utilities).
- Source Code Control systems (e.g. CVS, SVN).

ROS also has the ability to accept partial configuration updates. It is possible to, for example, update only the parameters for a single Ethernet port. Transferring a file containing only the following lines to a ROS device will result in an update of the parameters for Ethernet port 1 without changing any other parameters of the device’s configuration:

```
# Port Parameters
ethPortCfg
Port,Name,Media,State,AutoN,Speed,Dupx,FlowCtrl,LFI,Alarm,
1,Port 1,100TX,Enabled,On,Auto,Auto,Off,Off,On,
```

Applying the Configuration Update

Once a configuration file has been successfully transferred to a ROS device, irrespective of the transfer method, the device will reset itself automatically. Note that this behavior differs from that when upgrading firmware files, where a reset command must be issued by the administrator.

Security Considerations

The same limitations apply to writing **config.csv** to the ROS device that apply to firmware images. Refer to section Section 14.1, “ Upgrading Firmware ” for details on the permissions necessary to write the ROS configuration file.

14.3. Backing Up ROS System Files

All of the same file transfer mechanisms discussed in the preceding sections may also be used to transfer files **from** a ROS device, as well as to update firmware or configuration files. It might

be desirable, in addition to creating an archive of the device's firmware files, to back up the configuration database, **config.csv**, or system log file, **syslog.txt**, on a regular basis. Type "dir" at the ROS CLI for a listing and description of files on the ROS device.

An example of backing up a file using SFTP follows. For descriptions on the use of the other file transfer mechanisms, please refer to the examples in Section 14.1, " Upgrading Firmware ". Note that only the direction of file transfer changes.

14.3.1. Backing Up Files Using SFTP

This method requires that SFTP client software be available on a computer with a network connection to the ROS device that one wishes to back up. Establish an SFTP connection with administrative privileges to the ROS device. Begin transferring the desired file from the device. An example of using an SFTP session to create a local backup of the ROS main firmware image to a Linux workstation follows:

```
user31host$ sftp admin31ros_ip
Connecting to ros_ip...
admin31ros_ip's password:
sftp> get main.bin
Downloading /main.bin
main.bin                                      100% 2139KB   48.7KB/s   00:44
sftp>
```

All files in ROS may be backed up using an SFTP session with administrative privileges.

14.4. Using SQL Commands

The ROS provides an "SQL-like" command facility that allows expert users to perform several operations not possible under the user interface, namely:

- Restoring the contents of a specific table, but not the whole configuration, to their factory defaults.
- Search tables in the database for specific configurations.
- Make changes to tables predicated upon existing configurations.

When combined with RSH, SQL commands provide a means to query and configure large numbers of devices from a central location.

14.4.1. Getting Started

SQL information is obtainable via the CLI shell "SQL" command:

```
>sql

The SQL command provides an 'sql like' interface for manipulating all system
configuration and status parameters. Entering 'SQL HELP command-name' displays
detailed help for a specific command. Commands, clauses, table, and column
names are all case insensitive.
DEFAULT   Sets all records in a table(s) to factory defaults.
DELETE    Allows for records to be deleted from a table.
HELP      Provides help for any SQL command or clause.
INFO      Displays a variety of information about the tables in the database
```

```
INSERT    Allows for new records to be inserted into a table.
SAVE      Saves the database to non-volatile memory storage.
SELECT    Queries the database and displays selected records.
UPDATE    Allows for existing records in a table to be updated.
```

14.4.2. Finding the Correct Table

Many SQL commands operate upon specific tables in the database, and require the table name to be specified. Navigating the menu system to the desired menu and pressing <Ctrl-Z> will show the name of the table. The menu name and the corresponding database table name will be cited.

Another way to find a table name is to run the “sql info tables” command. This command also displays menu names and their corresponding database table names depending upon the features supported by the device:

```
Table                    Description
-------------------------------------------------------------------------
alarms                   Alarms
cpuDiags                 CPU Diagnostics
ethPortCfg               Port Parameters
ethPortStats             Ethernet Statistics
ethPortStatus            Port Status
ipCfg                    IP Services
```

14.4.3. Retrieving Information

Retrieving a Table

The SQL select subcommand is used to retrieve table information. The command, “sql select from ‘tablename’”, provides a summary of the parameters within the table, as well as their values:

```
>sql select from ipcfg

IP Address Type IP Address      Subnet          Gateway         Management VLAN
Inactivity Timeout Telnet Sessions Allowed Web Server Users Allowed TFTP Server
ModBus Address SSH Sessions Allowed
Static          10.90.0.2     255.0.0.0                     1
Disabled           8                       16                       Get Only
Disabled       8

1 records selected
```

Retrieving a Parameter from a Table

SQL select command may be used to retrieve a particular parameter from a table. SQL command “sql select parameter_name from tablename” is used for this purpose. The parameter name is always the same as those displayed in the menu system. If the parameter name has spaces in it (e.g. “IP Address”) the spaces must be replaced with underscores or the name must be quoted:

```
>sql select "ip address" from ipcfg

IP Address
192.168.0.8

1 records selected
```

Retrieving a Table with the ‘Where’ Clause

It is useful to be able to display specific rows of a table predicated upon the row having parameters of a specific value. Addition of “where” clause to the "select" statement will limit the results returned. For example, suppose that it is desirable to identify all ports on the device operating in Auto Select mode:

```
>sql select from ethportcfg where Media_Type = Auto_Select

Port Name            Status   Media Type  Flow Control FEFI     Link Alarms
5    Port 7          Enabled  Auto Select Enabled      Disabled Enabled
6    Port 8          Enabled  Auto Select Disabled     Disabled Enabled

2 records selected
```

It is also possible to select rows based on multiple parameters using "and" and "or" operations between comparisons in the "where" clause. For example:

```
>sql select from ethportcfg where Media_Type = Auto_Select and Flow_control =
 Disabled

Port Name            Status   Media Type  Flow Control FEFI     Link Alarms
6    Port 8          Enabled  Auto Select Disabled     Disabled Enabled

1 records selected
```

14.4.4. Changing Values in a Table

The “where” clause can be used to select rows in a table and to modify the fields in that row. As an example, suppose that it is desirable to identify all ports on the device operating in 100 Mbps full-duplex mode with flow control disabled, and to enable flow control on these ports:

```
>sql update ethportcfg set flow_control=enabled where ( media_type = Auto_Select
 and flow_control = disabled )

1 records updated
```

14.4.5. Setting Default Values in a Table

It is sometimes desirable to restore one table to its factory defaults without modifying the remainder of the configuration. The “sql default” command allows an individual table to be defaulted.

```
>sql default into ethportcfg
```

14.4.6. Using RSH and SQL

The combination of remote shell scripting and SQL commands offers a means to interrogate and maintain a large number of devices. Consistency of configuration across sites may be verified by this method. The following presents a simple example where the devices to interrogate are drawn from the file “Devices”:

```
C:> type Devices
10.0.1.1
10.0.1.2
10.0.1.3
```

```
c:\> for /F %i in (devices) do rsh %i -l admin,admin sql select from ethportcfg
 where flow_control = disabled

C:\>rsh 10.0.1.1 -l admin,admin sql select from ethportcfg where flow_control =
 disabled

Port Name            Status   Media Type  Flow Control FEFI     Link Alarms
5    Port 5          Enabled  Auto Select Disabled     Disabled Enabled

1 records selected

C:\>rsh 10.0.1.2 -l admin,admin sql select from ethportcfg where flow_control =
 disabled

0 records selected

C:\>rsh 10.0.1.3 -l admin,admin sql select from ethportcfg where flow_control =
 disabled

Port Name            Status   Media Type  Flow Control FEFI     Link Alarms
3    Port 3          Enabled  Auto Select Disabled     Disabled Enabled
7    Port 7          Enabled  Auto Select Disabled     Disabled Enabled
8    Port 8          Enabled  Auto Select Disabled     Disabled Enabled
13   Port 13         Enabled  Auto Select Disabled     Disabled Enabled

4 records selected

C:\
```

Appendix A. SNMP MIB Support

A.1. Standard MIBs

RFC	MODULE Name	Groups Supported
RFC 1907	SNMPv2-MIB	SNMP Group SNMP Community Group SNMP Set Group System Group SNMP Basic Notifications Group
RFC 2863	IF-MIB	General Information Group VHC Packet Group Counter Discontinuity Group Link Up/Down Notification Group
RFC 2011	IP-MIB	IP Group ICMP Group
RFC 2012	TCP-MIB	TCP Group
RFC 2013	UDP-MIB	UDP Group
RFC 2819	RMON-MIB	Ethernet Statistics Group History Groups (History Control Group and Ethernet History Group) Alarm Group Event Group
RFC 4188	BRIDGE-MIB	Base Bridge Group Base Port Group STP Bridge Group STP Port Group TP Bridge Group TP FDB Group TP Group Notification Group
RFC 4318	RSTP-MIB	Bridge Group Port Group
RFC 4363	Q-BRIDGE-MIB	qBridge Base Group qBridge Fdb Unicast Group qBridge Fdb Multicast Group

RFC	MODULE Name	Groups Supported
		qBridge VLAN Group qBridge VLAN Static Group qBridge Port Group
LLDP MIB	LLDP-MIB	LLDP Config Group LLDP Config Rx Group LLDP Config Tx Group LLDP Stats Rx Group LLDP Stats Tx Group LLDP Local System Group LLDP Remote System Group LLDP Notifications Group
RFC 3414	SNMP-USER-BASED-SM-MIB	Basic Group
RFC 3415	SNMP-VIEW-BASED-ACM-MIB	Basic Group

A.2. RuggedCom proprietary MIBs

Proprietary MIB	MODULE Name	Groups Supported
RuggedSwitch RuggedServer RuggedMC30	RUGGEDCOM-SWITCH-MIB RUGGEDCOM-SERVER-MIB RUGGEDCOM-MC30-MIB	Defines Agent Capabilities for Ruggedcom Switches Defines Agent Capabilities for Ruggedcom Servers Defines Agent Capabilities for RMC30
RuggedcomTraps	RUGGEDCOM-TRAPS-MIB	Generic Traps Group Power Supply Trap Group Notifications Group
RcSysInfo	RUGGEDCOM-SYS-INFO-MIB	System Error Objects Group, System Status Objects Group, System Objects Temperature Group, System Status Power Supply Group, System Info Device Info Group

Appendix B. SNMP Trap Summary

Appendix B – SNMP Trap Summary

The switch generates the standard traps summarized in the following table.

- from IF-MIB: linkDown, linkUp
- from SNMPv2-MIB: authenticationFailure coldStart
- from BRIDGE-MIB: newRoot, topologyChage
- from RMON-MIB: risingAlarm, fallingAlarm
- from LLDP-MIB: lldpRemoteTablesChange

The switch also generates the proprietary traps which are summarized in this document with their severity levels. These traps are described in the RC-TRAPS-MIB.

- genericTrap,
- powerSupplyTrap,
- swUpgradeTrap,
- cfgChangeTrap

Generic Traps carry information about event in severity and description objects. They are sent at the time an alarm is generated for the device. The following are examples of RuggedCom Generic Traps, along with the severity of each in brackets:

- heap error (alert)
- NTP server failure (notification)
- real time clock failure (error)
- failed password (warning)
- MAC address not learned by switch fabric (error)
- BootP client: TFTP transfer failure (error)
- received looped back BPDU (error)
- received two consecutive confusing BPDUs on port, forcing down (error)
- GVRP failed to learn – too many VLANs (warning)

Appendix C. List of Objects Eligible for RMON Alarms

The following table lists ROS® database objects which are eligible for RMON alarms:

snmpInPkts	The number of messages delivered to the SNMP Agent.
snmpInBadVersions	The total number of SNMP messages which were delivered to the SNMP Agent and were for an unsupported SNMP version.
snmpInBadCommunityNames	The total number of SNMP messages delivered to the SNMP Agent which used a unknown SNMP community name.
snmpInBadCommunityNames	The total number of SNMP messages delivered to the SNMP Agent which represented an SNMP operation which was not allowed by the SNMP community named in the message.
snmpInASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP Agent decoding received SNMP messages.
tcpActiveOpens	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state.
tcpPassiveOpens	The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state.
tcpAttemptFails	The number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT or the SYN-RCVD, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD.
tcpEstabResets	The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state
tcpCurrEstab	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE- WAIT.
tcpInSegs	The total number of segments received, including those received in error.
tcpOutSegs	The total number of segments sent, including those on current connections but excluding those containing only retransmitted bytes.
tcpRetransSegsDescr	The total number of segments retransmitted - that is, the number of TCP segments transmitted containing one or more previously transmitted bytes.
udpInDatagrams	The total number of UDP datagrams received and delivered to UDP users.
udpNoPorts	The total number of received UDP datagrams for which there was no application at the destination port.
udpInErrors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
udpOutDatagrams	The number of sent UDP datagrams.
rcDeviceStsTemperature	The temperature measured in the device.
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.

ipInAddrErrors	The number of input datagrams discarded because the IP address in their header's destination field was not a valid address to be received at this entity. This count includes invalid addresses and addresses of unsupported Classes. For entities which are not IP routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
ipOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this `no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default routers are down.
ipForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP routers, this counter will include only those packets which were Source-Routed via this entity, and the Source-route option processing was successful.
ipInUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
ipInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting reassembly
ipInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP)
ipOutRequests	The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.
ipOutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
ipFragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
ipFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
ipFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, e.g., because their Don't Fragment flag was set.
ipReasmFails	The number of IP datagrams successfully reassembled.
ipRasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
ifInOctets	The total number of bytes received on the interface, including framing characters.
ifOutOctets	The total number of bytes transmitted out of the interface.
ifInDiscards	The number of received packets that are dropped due to lack of receive buffers.
ifInUcastPkts	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which, were not addressed to a multicast or broadcast address at this sub-layer.

ifOutUcastPkts	The total number of transmitted packets which were not addressed to a multicast or broadcast address.
ifInNUcastPkts	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which, were addressed to a multicast or broadcast address at this sub-layer.
ifOutNUcastPkts	The total number of transmitted packets which were addressed to a multicast or broadcast address.
ifInErrors	The number of received packets that contained errors preventing them from being deliverable to a higher-layer protocol.
ifInMulticastPkts	The total number of good packets received that were directed to multicast address.
ifOutMulticastPkts	The total number of packets transmitted that were directed to multicast address.
ifInBroadcastPkts	The total number of good packets received that were directed to the broadcast address.
ifOutBroadcastPkts	The total number of packets transmitted that were directed to the broadcast address.
ifHCInUcastPkts	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which, were not addressed to a multicast or broadcast address at this sub-layer. This object is a 64 bit version of ifInUcastPkts.
ifOutUcastPkts	The total number of transmitted packets which were not addressed to a multicast or broadcast address. This object is a 64 bit version of ifOutUcastPkts.
ifHCInBroadcastPkts	The total number of good packets received that were directed to the broadcast address. This object is a 64 bit version of ifInBroadcastPkts.
ifHCOutBroadcastPkts	The total number of packets transmitted that were directed to the broadcast address. This object is a 64 bit version of ifOutBroadcastPkts.
ifHCInOctets	The total number of bytes received on the interface, including framing characters. This object is a 64 bit version of ifInOctets.
ifHCOutOctets	The total number of bytes transmitted out of the interface. This object is a 64 bit version of ifOutOctets.
ifHCInMulticastPkts	The total number of good packets received that were directed to multicast address.
ifHCOutMulticastPkts	The total number of packets transmitted that were directed to multicast address. This object is a 64 bit version of ifOutMulticastPkts.
dot1dBasePortMtuExceededDiscards	The number of frames discarded by this port due to an excessive size.
dot1dTpPortInFrames	The number of frames that have been received by this port from its segment.
dot1dTpPortOutFrames	The number of frames that have been transmitted by this port to its segment.
etherStatsDropEvents	The number of received packets that are dropped due to lack of receive buffers.
etherStatsOctets	The number of bytes in received good packets (Unicast+Multicast +Broadcast) and dropped packets.
etherStatsPkts	The number of received good packets (Unicast+Multicast+Broadcast) and dropped packets
etherStatsBroadcastPkts	The number of good Broadcast packets received.
etherStatsMulticastPkts	The number of good Multicast packets received.

etherStatsCRCAlignErrors	The number of packets received which meet all the following conditions:1. Packet data length is between 64 and 1536 bytes inclusive.2. Packet has invalid CRC.3. Collision Event has not been detected.4. Late Collision Event has not been detected.
etherStatsUndersizePkts	The number of received packets which meet all the following conditions:1. Packet data length is less than 64 bytes.2. Collision Event has not been detected.3. Late Collision Event has not been detected.4. Packet has valid CRC.
etherStatsOversizePkts	The number of packets received with data length greater than 1536 bytes and valid CRC.
etherStatsFragments	The number of packets received which meet all the following conditions:1. Packet data length is less than 642. Collision Event has not been detected.3. Late Collision Event has not been detected.4. CRC invalid.
etherStatsJabbers	The total number of packets received that were longer than 1518 bytes and had either a bad Frame Check Sequence or Alignment Error.
etherStatsCollisions	The best estimate of the total number of collisions on this Ethernet segment.
etherStatsPkts64Octets	The total number of received packets that were 64 bytes long.
etherStatsPkts65to127Octets	The total number of received packets that were between 65 and 127 bytes long.
etherStatsPkts128to255Octets	The total number of received packets that were between 128 and 255 bytes long.
etherStatsPkts256to511Octets	The total number of received packets that were between 256 and 511 bytes long.
etherStatsPkts512to1023Octets	The total number of received packets that were between 512 and 1023 bytes long.
etherStatsPkts1024to1518Octets	The total number of received packets that were between 1024 and 1518 bytes long.
lldpStatsRxPortFramesDiscardedTotal	The number of LLDP frames received by this LLDP agent on the indicated port and then discarded for any reason. This counter can provide an indication that LLDP header formatting problems may exist with the local LLDP agent in the sending system or that LLDPDU validation problems may exist with the local LLDP agent in the receiving system.
lldpStatsRxPortFramesErrors	The number of invalid LLDP frames received by this LLDP agent on the indicated port, while this LLDP agent is enabled.
lldpStatsRxPortFramesTotal	The number of valid LLDP frames received by this LLDP agent on the indicated port, while this LLDP agent is enabled.
lldpStatsRxPortTLVsDiscardedTotal	The number of LLDP TLVs discarded for any reason by this LLDP agent on the indicated port.
lldpStatsRxPortTLVsUnrecognizedTotal	The number of LLDP TLVs received on the given port that are not recognized by this LLDP agent on the indicated port.
lldpStatsRxPortAgeoutsTotal	The counter that represents the number of age-outs that occurred on a given port. An age-out is the number of times the complete set of information advertised by a neighbour has been deleted from tables contained in lldpRemoteSystemsData objects because the information timeliness interval has expired.
lldpStatsRemTablesInserts	The number of times the complete set of information has been inserted into tables contained in lldpRemoteSystemsData.

lldpStatsRemTablesDeletes	The number of times the complete set of information has been deleted from tables contained in lldpRemoteSystemsData objects.
lldpStatsRemTablesDrops	The number of times the complete set of information could not be entered into tables contained in lldpRemoteSystemsData objects because of insufficient resources.
lldpStatsRemTablesAgeouts	The number of times the complete set of information has been deleted from tables contained in lldpRemoteSystemsData objects because the information timeliness interval has expired.
dot1qVlanNumDeletes	The number of times a VLAN entry has been deleted from the dot1qVlanCurrentTable (for any reason). If an entry is deleted, then inserted, and then deleted, this counter will be incremented by 2.

Appendix D. ModBus Management Support and Memory Map

ModBus management support in RuggedCom devices provides the user with a simple interface with basic status information. Support for this protocol simplifies the job of SCADA System integrators who can now easily use this feature to retrieve basic info from RuggedCom devices via a familiar protocol. Predominantly read-only status information is provided though a few writable registers exist to provide operator commands.

The PDU format defined for the ModBus protocol is

Function Code	Data

The following ModBus function codes are supported by RuggedCom for device management through ModBus:

(1) Read Input Registers or Read Holding Registers – 0x04 or 0x03 for which the Modbus PDU would look like:

Request

Function code	1 Byte	0x04(0x03)
Starting Address	2 Bytes	0x0000 to 0xFFFF
Number of Input Registers	2 Bytes	0x0001 to 0x007D

Response

Function code	1 Byte	0x04(0x03)
Byte Count	1 Byte	2 x **N***
Input Registers	**N***X2 Bytes	

***N** = Quantity of Input Registers

(2) Write Multiple Registers – 0x10

Request

Function code	1 Byte	0x10
Starting Address	2 Bytes	0x0000 to 0xFFFF
Number of Registers	2 Bytes	0x0001 to 0x0079
Byte Count	1 Byte	2 x **N***
Registers Value	**N*** x 2 Bytes	Value of the register

***N** = Quantity of Input Registers

Response

Function code	1 Byte	0x10
Starting Address	2 Bytes	0x0000 to 0xFFFF
Number of Registers	2 Bytes	1 to 121 (0x79)

Note that, as RuggedCom devices have a variable number of ports, not all registers and bits apply to all products.

Registers that are not applicable to a given product return zero value.

E.g. registers referring to serial ports are not applicable to RuggedSwitch® products.

D.1. Modbus Memory Map

Address	#Registers	Description (Reference Table in UI)	R/W	Format
PRODUCT INFO (table Name: ProductInfo)				
0000	16	Product Identification	R	Text
0010	32	Firmware Identification	R	Text
0040	1	Number of Ethernet Ports	R	Uint16
0041	1	Number of Serial Ports	R	Uint16
0042	1	Number of Alarms	R	Uint16
0043	1	Power Supply Status	R	PSStatusCmd
PRODUCT WRITE REGISTERS (table Name: various tables)				
0080	1	Clear Alarms	W	Cmd
0081	2	Reset Ethernet Ports	W	PortCmd
0083	2	Clear Ethernet Statistics	W	PortCmd
0085	2	Reset Serial Ports	W	PortCmd
0087	2	Clear Serial Port Statistics	W	PortCmd
ALARMS (table Name: alarms)				
0100	64	Alarm 1	R	Alarm
0140	64	Alarm 2	R	Alarm
0180	64	Alarm 3	R	Alarm
01C0	64	Alarm 4	R	Alarm
0200	64	Alarm 5	R	Alarm
0240	64	Alarm 6	R	Alarm
0280	64	Alarm 7	R	Alarm
02C0	64	Alarm 8	R	Alarm
ETHERNET PORT STATUS (table Name: ethPortStats)				
03FE	2	Port Link Status	R	PortCmd
ETHERNET STATISTICS (table Name: rmonStats)				
0400	2	Port 1 Statistics - Ethernet In Packets	R	Uint32
0402	2	Port 2 Statistics - Ethernet In Packets	R	Uint32
0404	2	Port 3 Statistics - Ethernet In Packets	R	Uint32
0406	2	Port 4 Statistics - Ethernet In Packets	R	Uint32

Address	#Registers	Description (Reference Table in UI)	R/W	Format
0408	2	Port 5 Statistics - Ethernet In Packets	R	Uint32
040A	2	Port 6 Statistics - Ethernet In Packets	R	Uint32
040C	2	Port 7 Statistics - Ethernet In Packets	R	Uint32
040E	2	Port 8 Statistics - Ethernet In Packets	R	Uint32
0410	2	Port 9 Statistics - Ethernet In Packets	R	Uint32
0412	2	Port 10 Statistics - Ethernet In Packets	R	Uint32
0414	2	Port 11 Statistics - Ethernet In Packets	R	Uint32
0416	2	Port 12 Statistics - Ethernet In Packets	R	Uint32
0418	2	Port 13 Statistics - Ethernet In Packets	R	Uint32
041A	2	Port 14 Statistics - Ethernet In Packets	R	Uint32
041C	2	Port 15 Statistics - Ethernet In Packets	R	Uint32
041E	2	Port 16 Statistics - Ethernet In Packets	R	Uint32
Address	#Registers	Description	R/W	Format
0420	2	Port 17 Statistics - Ethernet In Packets	R	Uint32
0422	2	Port 18 Statistics - Ethernet In Packets	R	Uint32
0424	2	Port 19 Statistics - Ethernet In Packets	R	Uint32
0426	2	Port 20 Statistics - Ethernet In Packets	R	Uint32
0440	2	Port 1 Statistics - Ethernet Out Packets	R	Uint32
0442	2	Port 2 Statistics - Ethernet Out Packets	R	Uint32
0444	2	Port 3 Statistics - Ethernet Out Packets	R	Uint32
0446	2	Port 4 Statistics - Ethernet Out Packets	R	Uint32
0448	2	Port 5 Statistics - Ethernet Out Packets	R	Uint32
044A	2	Port 6 Statistics - Ethernet Out Packets	R	Uint32
044C	2	Port 7 Statistics - Ethernet Out Packets	R	Uint32
044E	2	Port 8 Statistics - Ethernet Out Packets	R	Uint32
0450	2	Port 9 Statistics - Ethernet Out Packets	R	Uint32
0452	2	Port 10 Statistics - Ethernet Out Packets	R	Uint32
0454	2	Port 11 Statistics - Ethernet Out Packets	R	Uint32
0456	2	Port 12 Statistics - Ethernet Out Packets	R	Uint32
0458	2	Port 13 Statistics - Ethernet Out Packets	R	Uint32
045A	2	Port 14 Statistics - Ethernet Out Packets	R	Uint32
045C	2	Port 15 Statistics - Ethernet Out Packets	R	Uint32
045E	2	Port 16 Statistics - Ethernet Out Packets	R	Uint32
0460	2	Port 17 Statistics - Ethernet Out Packets	R	Uint32
0462	2	Port 18 Statistics - Ethernet Out Packets	R	Uint32
0464	2	Port 19 Statistics - Ethernet Out Packets	R	Uint32
0466	2	Port 20 Statistics - Ethernet Out Packets	R	Uint32
0480	2	Port 1 Statistics - Ethernet In Octets	R	Uint32
0482	2	Port 2 Statistics - Ethernet In Octets	R	Uint32

Address	#Registers	Description (Reference Table in UI)	R/W	Format
0484	2	Port 3 Statistics - Ethernet In Octets	R	Uint32
0486	2	Port 4 Statistics - Ethernet In Octets	R	Uint32
0488	2	Port 5 Statistics - Ethernet In Octets	R	Uint32
048A	2	Port 6 Statistics - Ethernet In Octets	R	Uint32
048C	2	Port 7 Statistics - Ethernet In Octets	R	Uint32
048E	2	Port 8 Statistics - Ethernet In Octets	R	Uint32
0490	2	Port 9 Statistics - Ethernet In Octets	R	Uint32
0492	2	Port 10 Statistics - Ethernet In Octets	R	Uint32
0494	2	Port 11 Statistics - Ethernet In Octets	R	Uint32
0496	2	Port 12 Statistics - Ethernet In Octets	R	Uint32
0498	2	Port 13 Statistics - Ethernet In Octets	R	Uint32
049A	2	Port 14 Statistics - Ethernet In Octets	R	Uint32
049C	2	Port 15 Statistics - Ethernet In Octets	R	Uint32
049E	2	Port 16 Statistics - Ethernet In Octets	R	Uint32
04A0	2	Port 17 Statistics - Ethernet In Octets	R	Uint32
04A2	2	Port 18 Statistics - Ethernet In Octets	R	Uint32
04A4	2	Port 19 Statistics - Ethernet In Octets	R	Uint32
04A6	2	Port 20 Statistics - Ethernet In Octets	R	Uint32
04C0	2	Port 1 Statistics - Ethernet Out Octets	R	Uint32
Address	#Registers	Description	R/W	Format
04C2	2	Port 2 Statistics - Ethernet Out Octets	R	Uint32
04C4	2	Port 3 Statistics - Ethernet Out Octets	R	Uint32
04C6	2	Port 4 Statistics - Ethernet Out Octets	R	Uint32
04C8	2	Port 5 Statistics - Ethernet Out Octets	R	Uint32
04CA	2	Port 6 Statistics - Ethernet Out Octets	R	Uint32
04CC	2	Port 7 Statistics - Ethernet Out Octets	R	Uint32
04CE	2	Port 8 Statistics - Ethernet Out Octets	R	Uint32
04D0	2	Port 9 Statistics - Ethernet Out Octets	R	Uint32
04D2	2	Port 10 Statistics - Ethernet Out Octets	R	Uint32
04D4	2	Port 11 Statistics - Ethernet Out Octets	R	Uint32
04D6	2	Port 12 Statistics - Ethernet Out Octets	R	Uint32
04D8	2	Port 13 Statistics - Ethernet Out Octets	R	Uint32
04DA	2	Port 14 Statistics - Ethernet Out Octets	R	Uint32
04DC	2	Port 15 Statistics - Ethernet Out Octets	R	Uint32
04DE	2	Port 16 Statistics - Ethernet Out Octets	R	Uint32
04E0	2	Port 17 Statistics - Ethernet Out Octets	R	Uint32
04E2	2	Port 18 Statistics - Ethernet Out Octets	R	Uint32
04E4	2	Port 19 Statistics - Ethernet Out Octets	R	Uint32
04E6	2	Port 20 Statistics - Ethernet Out Octets	R	Uint32

Address	#Registers	Description (Reference Table in UI)	R/W	Format
SERIAL STATISTICS (table Name: uartPortStatus)				
0600	2	Port 1 Statistics – Serial In characters	R	Uint32
0602	2	Port 2 Statistics – Serial In characters	R	Uint32
0604	2	Port 3 Statistics – Serial In characters	R	Uint32
0606	2	Port 4 Statistics – Serial In characters	R	Uint32
0640	2	Port 1 Statistics – Serial Out characters	R	Uint32
0642	2	Port 2 Statistics – Serial Out characters	R	Uint32
0644	2	Port 3 Statistics – Serial Out characters	R	Uint32
0646	2	Port 4 Statistics – Serial Out characters	R	Uint32
0680	2	Port 1 Statistics – Serial In Packets	R	Uint32
0682	2	Port 2 Statistics – Serial In Packets	R	Uint32
0684	2	Port 3 Statistics – Serial In Packets	R	Uint32
0686	2	Port 4 Statistics – Serial In Packets	R	Uint32
06C0	2	Port 1 Statistics – Serial Out Packets	R	Uint32
06C2	2	Port 2 Statistics – Serial Out Packets	R	Uint32
06C4	2	Port 3 Statistics – Serial Out Packets	R	Uint32
06C6	2	Port 4 Statistics – Serial Out Packets	R	Uint32

D.1.1. Text

Simple ASCII representation of the information related to the product. ASCII characters' most significant byte of register comes first.

E.g. Read Multiple Registers request to read Product Identification from location 0x0000.

0x04	0x00	0x00	0x00	0x08

Response may look like:

0x04	0x10	0x53	0x59	0x53	0x54	0x45	0x4D	0x20	0x4E	0x41	0x4D	0x45
0x00	0x00	0x00	0x00	0x00								

Byte 3 until the end of the response above contains the ASCII representation of the product identification, "SYSTEM NAME". Since the length of this field is smaller than eight registers, the rest of the field are filled with zeros.

D.1.2. Cmd

This format is used to instruct the device to set the output to either 'true' or 'false'. The most significant byte comes first.

FF 00 hex requests output to be True.

00 00 hex requests output to be False.

Any value other than the suggested values does not affect the requested operation.

E.g. Write Multiple Registers request to clear alarms in the device.

0x10	0x00	0x80	0x00	0x01	2	0xFF	0x00

FF 00 for register 00 80 would clear the system alarms

00 00 would not clear any alarms

Response may look like:

0x10	0x00	0x80	0x00	0x01

D.1.3. Uint16

Standard Modbus 16 bit register.

D.1.4. Uint32

Standard 2 Modbus 16 bit registers – First register would hold higher 16 bits of the value and seconds register would hold lower 16 bits from the 32 bit value.

D.1.5. PortCmd

Descriptive bit layout (1= Requested Action is True; 0 bit = Requested Action is false) per port.

PortCmd at this time provides a bit layout of max of 32 ports hence utilizing 2 Modbus registers.

First Modbus register corresponds to ports 1 – 16.

Second Modbus register corresponds to ports 17 – 32 for a particular action.

Bits that do not apply to a particular product are always set to zero. See example for details.

Bit Value 1 – Requested action is true (e.g. This particular port is “Up”).

Bit Value 0 – Requested action is false (e.g. The particular port is “Down”).

Read Data from device using PortCmd:

E.g. A Modbus Request to read multiple registers from location – 0x03FE.

0x04	0x03	0xFE	0x00	0x02

Response would depend on the device as on how many ports are available on the device

E.g. If Max number of ports on RuggedCom device to which you are talking to is 20

Response may look like:

0x04	0x04	0xF2	0x76	0x00	0x05

In the above response, bytes 3 and 4 refer to register 1 - i.e. register at location 0X03FE indicating port status of ports 1–16 and Byte 5 and 6 representing register 2 at location 0x03FF would refer to port status from 17-32, though in this case since device has only 20 ports so byte 6 would contain the status for ports 17-20 starting from right to left. Rest of the bits in register 2 corresponding to non-existing ports would be zero.

Performing write actions on the device using PortCmd:

Write multiple register request to clear Ethernet port statistics

0x10	0x00	0x83	0x00	0x01	2	0x55	0x76	0x00	0x50

Bit value 1 implies clear Ethernet statistics on a corresponding port. Bit value 0 corresponding to a port means do nothing.

Response may look like:

0x10	0x00	0x81	0x00	0x02

D.1.6. Alarm

This format is also another form of text description. This text corresponds to the alarm description from the table holding all the alarms. Similar to the ‘Text ’ format this format would also have ASCII representation of alarms. Please note that alarms are stacked in RuggedCom device in the sequence of their occurrences. So first alarm on the stack would be Alarm1, next latched alarm in the device is Alarm 2 and so on. User has capability of seeing first 8 alarms from the stack if they exist. Zero value is sent if an alarm does not exist.

D.1.7. PSStatusCmd

Descriptive bit layout for providing the status of available power supplies in the unit. Bits 0-4 of lower byte of the register are used for this purpose.

Bits 0-1: Power Supply 1 Status

Bits 2-3: Power supply 2 Status

Rest of the bits in the register do not provide any system status info at this time

Interpretation of the values:

01: Power Supply not present (1)

10: Power Supply is functional (2)

11: Power Supply is not functional (3)

Values used for presenting power supply status have been derived from RuggedCom specific MIB for SNMP.

Read Power Supply Status from device using PSStatusCmd:

E.g. A Modbus Request to read multiple registers from location – 0x0043

0x04	0x00	0x43	0x00	0x01

Response may look like:

0x04	0x02	0x00	0x0A

In the above response lower byte of the register shows status of power supplies. As per the response both power supplies in the unit are functional.

Appendix E. Command Line Listing

The following commands are available at the command line of ROS®-based RuggedSwitch® and RuggedServer™ devices:

alarms	Displays list of available alarms. Usage: alarms [all] all - display all alarm instances (default empty) - display one instance of each alarm type.
clearalarms	Clears all alarms
clearethstats	Clears Ethernet statistics for one or more port(s) clearethstats ports|'all' 'ports' - comma separated port numbers (e.g. '1,3-5,7') 'all' - all ports
clearlogs	Clears the system and crash logs
cls	Clears the screen
delay	Pause a specified number of milliseconds
dir	Prints file directory listing
echo	Echoes the specified message to the screen
exit	Terminate this command line session
help	help [command name] [command name] - Name of command for which to get help. If no command is specified, a list of all available commands is displayed along with a brief description of each one
ipconfig	Displays IP configuration
login	Login to the shell i.e. set the access level
logout	Logout of the shell
ping	Usage: ping {dest} [count] [timeout] dest Target IP address. count Number of echo requests to send; default is 4. timeout Timeout in milliseconds to wait for each reply; range is 2-5000, default is 300 milliseconds.
reset	Perform a 'hard' reset of the switch
resetport	Reset one or more Ethernet ports which may be useful for forcing re-negotiation of speed and duplex or in situations where the link partner has latched into an inappropriate state. RESETPORT ports|'all' 'ports' - comma separated port numbers (e.g. '1,3-5,7') 'all' - all ports will be reset

rmon	Displays names of RMON alarm eligible objects
route	Displays gateway configuration
spuriouscount	Displays User Spurious Interrupt count
sql	The SQL command provides an 'sql like' interface for manipulating all system configuration and status parameters. Entering 'SQL HELP command-name' displays detailed help for a specific command. Commands, clauses, table, and column names are all case insensitive. DEFAULT Sets all records in a table(s) to factory defaults. DELETE Allows for records to be deleted from a table. HELP Provides help for any SQL command or clause. INFO Displays a variety of information about the tables in the database INSERT Enables new records to be inserted into a table. SAVE Saves the database to non-volatile memory storage. SELECT Queries the database and displays selected records. UPDATE Enables existing records in a table to be updated.
sudp	Usage: sudp ip_addr ip_port size num_msgs time_interval. This command sends test UDP broadcast datagram via source IP port to dest IP port. Parameters are defined as follows: ip_addr = destination IP address ip_port = source/destination IP port size = number of bytes in UDP datagram (ignored if data are entered) num_msgs = number of datagrams within time interval time_interval = basic interval in ms to send num_msgs
telnet	Usage: telnet dest dest: Server's IP address. NOTE: <Ctrl-C> closes telnet session
tftp	Usage: tftp server cmd fsource fdest server: Remote TFTP server's IP address cmd: put (upload) or get (download) fsource: Source filename dest: Destination filename NOTE: <Ctrl-C> stops a tftp transfer.
trace	Starts event tracing. Run "trace ?" for more help.
type	Displays the contents of a text file. Enter 'dir' for a directory listing of files.

	type filename
version	Prints software versions
xmodem	xmodem direction filename direction: send - send file to client receive - receive file from client filename: Enter 'dir' for list of all filenames

Index

A

Alarms
- Active Alarms, 185
- Critical Failure Relay, 186
- Passive Alarms, 186
- Using Alarms, 185

C

CLI Shell Command
- clearlogs, 197
- dir, 196
- help, 196
- ipconfig, 200
- ping, 197
- reset, 201
- Summary, 195
- trace, 198
- type, 197

Configuration
- Update, 205

CoS
- Configuration, 147
- CoS (Classes of Service) Operation, 146
- DSCP, 152

CoS Configuration, 142

D

DHCP Relay, 45

Diagnostics
- CPU Diagnostics, 190
- Product Information, 192
- System Log, 191

E

Ethernet
- Cable Diagnostics, 58
- Link Detection, 62
- Loss-of-Link Management, 50
- Port Configuration, 50, 51
- Port Mirroring, 57
- Port Rate Limiting, 55
- Resetting Ports, 64
- Troubleshooting Ports, 64

F

Factory Default Configuration, Loading, 193

Firmware
- SFTP Upgrade, 205
- TFTP Client Upgrade, 204
- TFTP Server Upgrade, 204
- Upgrade, 202
- XModem Upgrade, 203

G

GMRP, 158
- Configuration, 166
- Example, 159
- Joining, 158
- Leaving, 159

I

IGMP, 154
- Active and Passive Mode, 155
- Configuration, 162
- Consumers and Producers, 154
- general membership query, 155
- group-specific membership query, 155
- leave group message, 155
- membership report, 155
- Operation, 154

Interface, 14

L

Layer 3 switches, Using, 125

LLDP
- Configuration, 176

M

MAC Addresses
- Configuring, 172
- Learning Options, 172
- Purging, 174
- Viewing, 171

MSTP
- Benefits, 98
- Boundary Port, 98
- CIST, 97
- CST, 97
- Implementing, 99
- IST, 97
- Master Port, 98
- MSTI, 96
- Operation, 95
- Region, 96

Multicast Filtering, 154

Configuration, 161
Static Configuration, 164

N

Network Discovery, 175
RCDP, 176

P

Passwords, 28

R

RADIUS, 40
Reset, Device, 193
RMON, 72
RMON Event Configuration, 82
ROS
RS232 Console Interface, 14
Secure Shell Server, 16
Web Server, 17
RSH, 200
RSTP
Bridge Diameter, 95
Bridge Parameters, 105
Edge Ports, 94
Operation, 91
Path Costs, 94
Point To Point Links, 94
Port Parameters, 107
Port Redundancy, 103
Ring Backbone Configurations, 101
Statistics, Bridge, 114
Statistics, Port, 116
Structured Wiring Configurations, 100
Troubleshooting, 122

S

SNMP Management, 35
SQL
Default Command, 209
Info Command, 208
Select Command, 208
Update Command, 209
Using Commands, 207
'From' Clause, 208
'Where' Clause, 209
Syslog, 46

T

TACACS+, 43
Time Synchronization, 31
Troubleshooting
Multicast Filtering, 167

V

VLAN
Configuration, 132
Configuring Static VLANs, 134
Displaying VLANs, 137
Edge Type, 126
GVRP, 128
Ingress and Egress Rules, 127
Management VLAN, 126
Native, 125
Operation, 125
Port Parameters, 135
PVLAN Edge, 129
QinQ, 130
Tagging, 125
Troubleshooting, 138
Trunk Type, 126